PCT WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 : H04M 3/50, 7/00, H04L 12/58, H04Q 3/00 — A1

(11) International Publication Number: **WO 98/32272**

(43) International Publication Date: 23 July 1998 (23.07.98)

(21) International Application Number: PCT/GB98/00126

(22) International Filing Date: 15 January 1998 (15.01.98)

(30) Priority Data:
08/784,131 15 January 1997 (15.01.97) US

(63) Related by Continuation (CON) or Continuation-in-Part (CIP) to Earlier Application
US 08/784,131 (CIP)
Filed on 15 January 1997 (15.01.97)

(71) Applicant *(for all designated States except US)*: BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY [GB/GB]; 81 Newgate Street, London EC1A 7AJ (GB).

(72) Inventors; and
(75) Inventors/Applicants *(for US only)*: LEGGATT, Jon, Stewart [GB/GB]; 12B Fonnereau Road, Ipswich, Suffolk IP1 3JP (GB). CAPP, Matthew, William [GB/GB]; 8 Scopes Road, Kesgrave, Ipswich, Suffolk IP3 9NU (GB). ASTIZ, Paul [US/US]; 11115 Whisperwood Lane, North Bethesda, MD 20852 (US).

(74) Agent: DUTTON, Erica, Lindley, Graham; BT Group Legal Services, Intellectual Property Dept., Holborn Centre, 8th floor, 120 Holborn, London EC1N 2TE (GB).

(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GM, GW, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).

Published
With international search report.
Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.

(54) Title: METHOD AND APPARATUS FOR MESSAGING

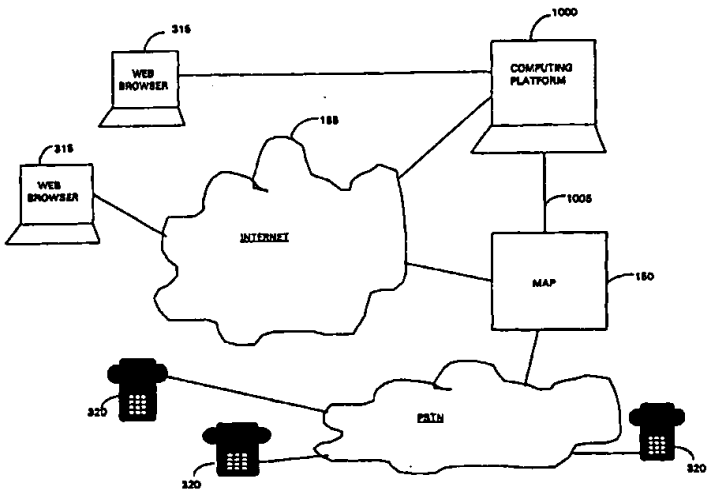

(57) Abstract

A Web server (1000) and Web browser (315) are used to structure message requests for a platform (150) which bridges the Internet (155) and a telecommunications network. The structured message requests include message content information and destination information. The destination information optionally can contain multiple destination identifiers. A single message request can then be sent to the platform (150) where it is used to send common message content to a plurality of destinations. The message requests may be sent in SMTP compatible format and the platform be provided with text to speech translation capability such that the platform can output the message content to users connected to a telecommunications network. The Web server has access to a customer-related data store so that the customer can generate a directory of destinations and can precompose messages. Further, the Web server has access to call record data and can download substantially real-time call status information stored as call record data by the platform (150).

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	**ES**	Spain	**LS**	Lesotho	**SI**	Slovenia
AM	Armenia	**FI**	Finland	**LT**	Lithuania	**SK**	Slovakia
AT	Austria	**FR**	France	**LU**	Luxembourg	**SN**	Senegal
AU	Australia	**GA**	Gabon	**LV**	Latvia	**SZ**	Swaziland
AZ	Azerbaijan	**GB**	United Kingdom	**MC**	Monaco	**TD**	Chad
BA	Bosnia and Herzegovina	**GE**	Georgia	**MD**	Republic of Moldova	**TG**	Togo
BB	Barbados	**GH**	Ghana	**MG**	Madagascar	**TJ**	Tajikistan
BE	Belgium	**GN**	Guinea	**MK**	The former Yugoslav Republic of Macedonia	**TM**	Turkmenistan
BF	Burkina Faso	**GR**	Greece			**TR**	Turkey
BG	Bulgaria	**HU**	Hungary	**ML**	Mali	**TT**	Trinidad and Tobago
BJ	Benin	**IE**	Ireland	**MN**	Mongolia	**UA**	Ukraine
BR	Brazil	**IL**	Israel	**MR**	Mauritania	**UG**	Uganda
BY	Belarus	**IS**	Iceland	**MW**	Malawi	**US**	United States of America
CA	Canada	**IT**	Italy	**MX**	Mexico	**UZ**	Uzbekistan
CF	Central African Republic	**JP**	Japan	**NE**	Niger	**VN**	Viet Nam
CG	Congo	**KE**	Kenya	**NL**	Netherlands	**YU**	Yugoslavia
CH	Switzerland	**KG**	Kyrgyzstan	**NO**	Norway	**ZW**	Zimbabwe
CI	Côte d'Ivoire	**KP**	Democratic People's Republic of Korea	**NZ**	New Zealand		
CM	Cameroon			**PL**	Poland		
CN	China	**KR**	Republic of Korea	**PT**	Portugal		
CU	Cuba	**KZ**	Kazakstan	**RO**	Romania		
CZ	Czech Republic	**LC**	Saint Lucia	**RU**	Russian Federation		
DE	Germany	**LI**	Liechtenstein	**SD**	Sudan		
DK	Denmark	**LK**	Sri Lanka	**SE**	Sweden		
EE	Estonia	**LR**	Liberia	**SG**	Singapore		

METHOD AND APPARATUS FOR MESSAGING

The present invention relates to method and apparatus for messaging in a communications network and finds particular application in an integrated communications environment providing different service types, such as voice telephony and messaging.

A system which can provide improved access to both telephony and other communication types, such as text messages and electronic mail, is described in copending British patent application number 9619958.3, filed on 25 September 1996 by the present applicant, the subject matter of which is herein incorporated by reference. This system is referred to below as the Minor Applications Platform (MAP). It is marketed by Aculab plc under the product name "Millenium CT".

The MAP essentially provides a bridge between different communication types. It can for instance receive service requests as SMTP inputs, or inputs over a telephone line, and respond by providing a facsimile or electronic mail (e/mail) service as an output.

According to the present invention, there is provided a messaging system for use in providing a messaging service, which system comprises:

i) an input for receiving a messaging service request, the service request comprising destination information and message content information;

ii) a user data interface and store for receiving and storing user input data;

iii) request processing means for processing a received service request; and

iv) an output for outputting a processed service request in a format for use by communications connection set-up means in transmitting the message content information, or information identified thereby, as an

output message to at least one destination identified by the destination information.

Embodiments of the present invention will need to have access to a communications connection set-up means capable of receiving and responding appropriately to processed service requests and this can advantageously be provided by a MAP, as abovementioned. Such a system might be as described above, but further comprising the connection set-up means, said connection set-up means comprising:

i) a voice network interface, for providing access to a network for carrying voice signals;

ii) a data network interface, for receiving processed service requests;

iii) a resource interface for access to resources for use in providing communications services, including at least one speech-related resource from the group comprising voice recognition, recordal of incoming sound signals and transmission of outgoing sound signals;

iv) interpretation means for use in relating a processed service request to a computer-based application for use in provision of that service; and

v) initiating means for initiating the running of a computer-based application related to a processed service request

so as to generate a common voice message output to more than one location in the network carrying voice signals.

Other connection set-up means could be used however. The basic requirement will be for platform, which may incorporate a combination of hardware and software, which is provided with an appropriate interface for receiving the processed service requests, and the means to respond appropriately by establishing a connection to a

location in a communications network and sending the message content information, or information identified thereby.

In embodiments of the present invention, a flexible messaging service can be provided which can make considerable management information and control functions available to the user.

The access to data storage for storing customer-related data gives considerable flexibility to the user. This is because the request processing means can use customer-related data to interpret a received service request. The destination information in a received service request need not then be in a form directly usable by the connection set-up means, instead being interpreted by the request processing means. For instance, the user could create a personal directory of destinations for output messages, listed under peoples' names and groups of names. The user could then use names in a service request. The names are interpreted by the request processing means into location identifiers for use by the connection set-up means, using the customer-related data.

The user could create address groups and use simply a group name in the initial service request, this group name then being interpreted into multiple location identifiers by the request processing means.

A location identifier in this context, which could be used directly by the connection set-up means, could be for instance a telephone number, an electronic mail address or a Universal Resource Locator (URL). That is, it generally will address a network location to which a communications connection can be established.

The request processing means could be provided with access to control data, in addition to or as part of the customer-related data. It could then be used as a mechanism for instance for screening or limiting destinations for the output messages.

As well as interpreting destination information to obtain location identifiers, the request processing means could also interpret the message content information by reference to the data storage. Hence a user could prepare stock messages, or a particularly long message, in advance and simply identify the message to be sent by referencing a message in a message directory.

The system can be installed on an existing platform, for instance on a computer which provides Web server capability and which is connected to the Internet. The system itself, however, may optionally be made only accessible from a corporate Intranet, or to another closed user group. By using a MAP which is connected for instance to both the Internet and a PSTN, messages can be sent to a very broad spectrum of destinations. On the other hand, the output could be limited in the same manner as user access, for instance by use of data in the data storage, to only predesignated destinations.

Embodiments of the present invention can be particularly efficient in terms of network traffic since the received service request comprises a single message to a single destination. However, the service request can potentially contain destination information to multiple locations for the same message content information ("multimessaging"). It need only be transformed into the multiple messages, the output messages, after it has reached the communications connection set-up means.

If an embodiment of the present invention was provided on computing platform having Windows capability, for instance providing a Web server in a company "Intranet", a convenient means by which a user may create a messaging service request is by means of a computing terminal provided with a forms facility, such as a Web browser. This can elicit the necessary data from the user which the request processing means converts to a message suitable for use by the communications connection set-up means.

In some environments, it may be preferable that the system has access to text-to-speech conversion means such that the message content information can be delivered as output voice messages. Where the service requests are received in text or data format, they can then be delivered to any person having access to a telephone and do not require the use of more expensive forms of equipment for reception. In environments in which other equipment is available to the expected recipients, however, such as in a corporate environment, the message content information could be delivered by a different means, such as in a textual format of the e/mail type.

A message output which can provide appropriate real-time communication from the messaging system to a MAP is by use of the UNIX-based development known as a "socket" connection. By using a direct, point-to-point connection, messages in e/mail format can be sent in real time to the MAP. This is further discussed below. It avoids delays experienced in normal e/mail transmissions while allowing user-friendly textual inputs to be simply formatted into a processed service request which the MAP is easily adapted to receive and act on.

The word "message" should not be taken to have a specific meaning in the context of the present specification, unless specifically stated as having such a meaning. In particular, the word "message" should not be taken to imply for instance any particular means of transmission, routing or queuing, as might be implied for instance by e/mail (i.e. electronic mail) messages. It is intended only to mean a set of information which can be transferred between network entities. A message could thus be embodied as speech, text or data signal.

Further, although the present specification refers to a "user" inputting service requests, and to "user data" and the like, the "user" could in practice be another software system or piece of equipment. For instance, a service request could be automatically generated by a metering device when a metered parameter goes below or above a predetermined threshold.

A messaging system and service, using a MAP for text to speech conversion and for delivering the output messages to their destinations, will now be described as a specific embodiment of the invention, by way of example only, with reference to the accompanying drawings in which:

Figure 1 shows a block diagram of the hardware context of the messaging system;

Figure 2 shows a block diagram of components of the system;

Figure 3 shows a flow diagram of a registration process for a user of the messaging system of Figure 2;

Figure 4 shows a flow diagram for a directory creation and modification process for a user of the messaging system of Figure2;

Figure 5 shows a flow diagram for a password process for a user of the messaging system of Figure2;

Figure 6 shows a flow diagram for a service request process available to a user of the messaging system of Figure 2;

Figure 7 shows shows a flow diagram for a call records process available to a user of the messaging system of Figure 2;

Figure 8 shows a schematic block diagram of components of a MAP for use with the messaging system of Figure 2;

Figure 9 shows a block diagram of the contents of a speech card for use in the MAP shown in Figure 8;

Figure 10 shows a schematic block diagram of two MAPs according to Figure 8, having some shared platform;

Figure 11 shows a functional block diagram of a MAP management system for managing overall process flow in the MAP of Figure 8;

Figure 12 shows a schematic representation of an example of data provision at startup for the MAP shown in Figure 8;

Figure 13 shows a flow diagram of events occurring in the MAP of Figure 8, at launch of an application in response to a service request; and

Figure 14 shows a schematic overview of service and resource queues in the MAP of Figure 8.

GLOSSARY OF ACRONYMS USED IN THIS SPECIFICATION

API	Applications Programming Interface
ASCII	American Standard Code for Information Interchange
ASIIWR	Automatic Speaker Independent Isolated Word Recogniser
BT	British Telecommunications plc (the present applicant)
CDHMM	Continuous Density Hidden Markov Model
CGI	Common Gateway Interface
CSS	Customer Support Systems
DASS2	Digital Automatic Signalling System No. 2 (used in UK with provision of primary rate ISDN)
DDI	Direct Dialled In
DLIC	Digital Line Interface Card
DPNSS	Digital Private Network Signalling System (generally used between PABXs, has core capabilities and manufacturer - specific extensions).
DSP	Digital Signal Processor
DTMF	Dual Tone Multi-Frequency
DX CPU	DX Central Processing Unit
EMC	Electromagnetic Compatibility
FTP	File Transfer Protocol
GUI	Graphical User Interface
HTML	HyperText Markup Language
IDE	
ISA	Industry Standard Architecture
ISDN	Integrated Services Digital Network
LAN	Local Area Network
MAP	Minor Applications Platform
MVIP	Multi-Vendor Interface Protocol
NIC	Network Interface Card
NT	Trade Mark
ODBC	Open Database Connectivity
PABX	Private Automatic Branch Exchange

PC	Personal Computer
PCM	Pulse Code Modulation
PIN	Personal Identification Number
PSTN	Public Switched Telecommunications Network
SC	Speech Card
SIC	Service Identity Code
SMS	Short Messaging Service
SMTP	Simple Mail Transfer Protocol
STAP	Speech Technology Applications Processor
STD	Standard Trunk Dialling (British area code convention)
T120	An ITU standard
TCP/IP	Transmission Control Protocol/Internet Protocol
TLI	Transfer Line Identity
TTS(M/F)	Text to Speech (Male/Female)
URL	Universal Resource Locator
VM	Voice Messaging
WWW (or Web)	World Wide Web

MESSAGING SYSTEM

Referring to Figure 1, the messaging system is provided on a computing platform 1000 provided with a Web server 1010 and connected via a data connection 1005 to a MAP 150. Web browsers 315 for multiple users are connected to the Web server 1010 and the MAP 150 is connected to a PSTN 160. Multiple users are provided with telephone connections 320 to the PSTN 160, in the normal way.

The Web browsers 315 are standard browsers, for instance as offered by Netscape, modified to conform to CGI capabilities by the use of the C++ programming language. A suitable Web server is for instance the Netscape "Commerce Server, Version 1.1".

The data connection 1005 between the Web Server 1010 and the MAP 150 carries TCP/IP data and is preferably provided by a real-time, point to point connection such as provided by the UNIX capability known as a "socket" connection. This allows fast communication with the MAP 150 rather than the performance of an ordinary electronic mail connection.

Socket connections are known and information can be obtained by reference to "The design of the UNIX Operating system", by Maurice J. Bach, Prentice Hall International, 1986. A further reference for information is the Unix Programmers Manual 4.2, Berkeley Software Distribution, Virtual Vax-11 version, obtainable from or published by the Computer Science Division, Dept of Electrical Engineering and Computer Science, University of California at Berkeley, August 1983.

The relationship between the computing platform 1000 which supports the messaging system and the MAP 150 is that the platform 1000 supports most of the processing capacity for the service and in particular provides user-related aspects of the service. That is, the platform 1000 provides the means for inputting data from the user, formatting output messages to the MAP 150 based on user input, creating and updating a user directory, creating and maintaining user account data and secure user access and maintaining user-specific call records and status data. The MAP 150 provides the functionality for converting a text message received from the messaging platform 1000 to speech, scheduling and making telephone calls using that speech, making retries and updating the call records and the status data available to the Web server 1010.

Referring to Figures 2 and 10, the computing platform 1000 supports the Web server 1010, an HTML page store 2045, data storage 2030, 2035 and a set of gateway processes 2000, 2005, 2010, 2015, 2020.

In general, the computing platform 1000 can be accessed by means of a user's terminal 315 and can access the MAP 150 in any of several different ways. For instance, it could have a socket connection to the MAP, a LAN connection, or it could be connected via the Internet 155. The messaging service software may be

installed on computing platform 1000 which comprises one or more personal computers (PCs) running WINDOWS '95, or an NT workstation. In addition it may be possible for the software to run on WINDOWS 3.1 or with Windows for Workgroups V3.11 where the WIN32S extensions from Microsoft will be required.

The data storage 2030, 2035 could be located elsewhere and merely be accessible via the platform 1000. Further, although shown in separate units the data storage could be provided as a single entity.

A user of the messaging system has access from a Web browser 315 which communicates with the Web server 1010 in known manner. The Web server 1010 provides Web pages to the browser 315 from the HTML page store 2045. More frequently however, the Web server 1010 will send pages generated by the gateway processes, using a common gateway interface (CGI) 2040. Hence an input from a user will usually identify a gateway process, by means of a URL, which the CGI can recognise and initiate in order to create a Web page for transmission back to the user and to respond to user inputs appropriately.

It might be noted that gateway processes do not normally maintain context once they have been run. There is a requirement in embodiments of the present invention for maintaining context between one instance of running a gateway process and another. This is done by embedding context data in the relevant URLs, using the known "GET" and "POST" techniques which pass data in the query string or message body of a URL.

An example is:

http://123.113.59.58/aservice/mms.cgi?acc=XxYZaB.20oFra&pg=home

This URL identifies the home page for a multimessaging service, with embedded context information in the form of a user account number. It has the following sections:

"123-113.59.58" is the machine IP number for the Web Server 1010

"aservice" is the directory in the machine that the multimessaging

	program(s) are in
"mms.cgi"	is the name of the multimessaging program
"XxYZaB.20oFr"	is the user account number, encrypted
"home"	names the specific IP page.

GATEWAY PROCESSES IN THE MESSAGING SYSTEM

There are five main gateway process types provided by the Web server 1010:

- Registration 2000
- User Directory 2005
- Password 2010
- Service Request 2015
- Call Records 2020

There will be other gateway processes, such as a login process, but these are of known type and functionality and are not shown or described herein.

From the point of view of a user at a Web browser 315, the registration process 2000 is seen separately from the other processes since it is only necessary to register for the messaging service once. Thereafter, the information provided/generated at registration usually remains unchanged. The other processes are offered as alternatives from a common Web page for the user to select from, since the user will normally use each of them many times.

Referring to Figure 3, the registration process is relatively straight forward.

STEP 3000: It is triggered after the user inputs a URL at the Web browser for the messaging service home page.

STEP 3005: In known manner, the Web server downloads the relevant home page to the Web browser.

STEP 3010: The home page is provided with a selection of HTML links, for instance as "Hotlinks", and the user selects registration.

STEP 3015: The CGI 2040 initiates the registration process and downloads a form to the user.

STEP 3020: The user completes the form, including provision of a password, and returns it to the server.

STEP 3025: The registration process performs a check on the completed form and either offers the user a retry (for a limited number of times) if the completed form fails to meet requirements (STEP 3035), or allocates an account number, which it notifies to the user, and stores the user data in association with the account number in the user data store 2035 with password only access (STEP 3030). The user data required at this stage is of course optional but might include the user's name and telephone number.

Subsequent to registration, the login procedure for a user will usually require both the account number and the password to be provided. This provides security and allows one user to run different accounts.

Referring to Figure 4, the user directory process 2005 is as follows:

STEPS 3000, 3005: As described with reference to Figure 3. (It is of course not necessary that the "Hotlinks" for all the processes from which a user can select are presented together on one Web page. There may be an intervening page or pages to be selected by the user before the page with a required "Hotlink" appears. This is simply standard HTML technology.)

STEP 4000: The user selects an HTML link identifying the user directory process 2005. (In practice, there are two separate "Hotlinks", one for adding to a personal directory and one for removing entries from a personal directory. The two processes are complementary and only the process of adding entries to a user directory is described here.)

STEP 4005: The CGI 2040 initiates the requested User Directory Process 2005 and returns a form to the user, requesting information concerning a new entry to the user directory. The information will usually comprise a name and telephone number (or other network location identifier) for the new entry. The form also

requires a user password. When the user selects the command "add", the process moves to the next step.

STEP 3025: The process checks whether the completed form meets requirements, including particularly the provision of a valid password, and if not notifies the user and offers a limited number of retries (STEP 3035), as in the process of Figure 3.

STEP 4010: When the completed form is found to meet requirements, the process enters the data in a directory allocated to the user account number in the gateway data store 2035. The process will also notify the user that the entry has been added, using a Web Page which offers a Hotlink to return to the service home page.

Referring to Figure 5, the password process 2010 is as follows:

STEPS 3000, 3005: As described with reference to Figure 3.

STEP 5000: The user selects an HTML link identifying the password process 2010.

STEP 5005: The CGI 2040 initiates the requested password process 2010 and returns a form to the user, requesting information concerning old and new passwords.

STEP 5010: The user returns a completed form.

STEP 3025: The process checks whether the completed form meets requirements and, if not, notifies the user and offers a limited number of retries (STEP 3035), as in the process of Figure 3.

STEP 5015: When the completed form is found to meet requirements, the process substitutes the new password for the old password in all relevant locations, including in a directory allocated to the user account number in the gateway data store 2035. The process will also notify the user that the password change has been accepted, using a Web Page which offers a Hotlink to return to the service home page.

Referring to Figure 6, the messaging service process 2015 is as follows:

STEPS 3000, 3005: As described with reference to Figure 3.

STEP 6000: The user selects an HTML link identifying the messaging service process 2015.

STEP 6005: The CGI 2040 initiates the requested Messaging Service Process 2010 and returns a form to the user.

STEP 6010: The user returns a completed form.

STEP 3025: The process checks whether the completed form meets requirements and, if not, notifies the user and offers a limited number of retries (STEP 3035), as in the process of Figure 3.

STEP 6015: When the completed form is found to meet requirements, the process uses the information supplied to generate an SMTP message for transmission to the MAP 150. The process will also notify the user that the message has been sent, identifying the message, for instance by its title, and giving time and date information. Notification is done by means of a Web Page which offers a Hotlink to return to the service home page.

The SMTP message generated at STEP 6015 is required to contain specific fields for data in order that the MAP can receive and respond to the message. The following gives a specification of each of the variables in the MAP service messages so as to clarify the protocol between a client application and the MAP.

Messages to the MAP will be addressed in the following form:

<user>@<service>.map.cartoon.bt.co.uk

The user field

The <user> field may contain the telephone number of the person requesting the service.

The service field

The <service> field contains the name of the service being requested. Some examples are given below:

ttsm	text to speech - male voice
ttsf	text to speech - female voice
tel	telephony
page	to send a page message (may need to split according to swatch, data pagers etc)
sms	sends an SMS (Sort Message Service) to a GSM phone
fax	fax message
conf	telephone conferencing

All services will exchange the same message format. Embodiments of the present invention which require a voice output message from the MAP 150 will generally use the name "ttsm" or "ttsf". Redundant fields will be assigned the string NULL. This makes it easier to code the client and server as well as making it easier to manage the interface.

The structure of messages incoming to the MAP is generally as set out below. However, it would be relatively simple to modify applications at the MAP 150 so that other message elements could be incorporated and acted on.

service_instance_identify =
- the identifier assigned to the service_instance by the MAP (initially this will be 0000)

start_date_time =
- the start date and time given as Mmddhhmm

stop_date_time =
- the stop date and time given as Mmddhhmm

sender_telephone_number =
- the telephone number of the person setting up the service_instance including STD code

sender_name =
- the real name of the person setting up the service_instance

sender_email_address =

- the email of the person setting up the service_instance

sender_fax_number =

- the fax number of the person setting up the service_instance

sender_pager_number =

- the pager number of the person setting up the service_instance

sender_SMS_number =

- the SMS number of the person setting up the service_instance (GSM phone number)

number_of_receivers =

- the number of people receiving the service_instance (not including the sender)

receivers_telephone_number =

- a list of the tel numbers of the receivers (separated by whitespace)

receivers_email_address =

- a list of the emails of the receivers (separated by whitespace)
- same order as receivers_telephone_number

receivers_fax_number =

- a list of the fax numbers of the receivers (separated by whitespace)
- same order as receivers_telephone_number

receivers_pager_number =

- a list of the pager numbers of the receivers (separated by whitespace)
- same order as receivers_telephone_number

receivers_SMS_number =

- a list of the SMS numbers of the receivers (separated by whitespace)
- same order as receivers_telephone_number

confirmation_email_address =

- the address to which the confirmation that the service_instance has been successful should be sent

start_now =

- set to YES if the service_instance is to be started as soon as possible
- set to NO if the service_instance is to be started at the specified time

delete_service_instance =

- set to YES if the service_instance is to be deleted

- set to NO if the service_instance is not to be deleted

message_start =

- flags the beginning of the message body

message_stop =

- flags the end of the message body

status =

- set to FAIL if the TTS service_instance has been rejected
- set to SUCCESS if the TTS service_instance has been accepted

Parts of the message structure set out above are further discussed below. (It might be noted that not all of the fields in the above message structure will be necessary in every embodiment of the present invention. The full message structure is a requirement for supporting additional services provided by the MAP.)

"Service_instance_identify" is an identifier assigned by the MAP on first receipt of a processed service request. It can be used thereafter wherever information will be relevant to that particular service instance. Hence it will be used by the messaging system to identify status data in the raw call records data store 2030 for downloading to the user so that the user can check whether their message has been sent or has failed.

"Start_date_time" and "stop_date_time" allow the user to schedule when a message should be sent. The MAP will only start to try to deliver a message after the start date and time given and will terminate attempts after the stop date and time given. The user can specify a time and date at which each message is to be delivered and this can be used for arranging meetings, scheduling events, providing public information and raising alarms.

The sender details which make up the next six data fields are not necessary for embodiments of the present invention which only require feedback from the MAP

in the form of raw call record data. They would therefore normally be assigned the string "NULL".

The "number_of_receivers" field could be completed by the service request process 2015 by processing data entered by the user. It could alternatively be filled in by the user. Either way, it can be used by the system to make a simple check that network location information, such as an e/mail address, has been supplied for the correct number of intended recipients of a message. If the check fails, the system will issue a prompt to the user for missing information.

The five fields for receivers' data will generally not all be used in the same message since usually the intention will be that messages are to be delivered by a selected one only of a set of possible modes of transmission. However, in more sophisticated embodiments of the present invention, it would be feasible to instruct the MAP to use an alternative mode of transmission if a first one should fail. In that case, receiver location information is likely to be needed for each of the alternative forms of delivery selected for the alternative modes.

"Confirmation_email_address" is a field which could potentially be used in embodiments of the present invention. However, in the primary embodiment described herein, this functionality of the MAP is not used. Instead the messaging system uses its own processing power to access raw call record data so as to monitor progress of a service instance.

"Start_now" will be set to NO in the event that the "start_date_time" field has an entry. Otherwise it will be set to YES.

"Delete_service_instance" is unlikely to be completed in a processed service request. However, it could clearly be completed by the messaging service in other circumstances, for instance in the event that a user had the option to cancel a service request before the MAP had acted on it.

"Message_start" and "message_stop" identify the information content that should be incorporated in a message output to a receiver by the MAP. For instance, in an example of functionality provided by the MAP to deliver a message according to an embodiment of the present invention, the MAP may set up a telephone connection to a receiving party. If the receiving party answers, the MAP will play a recorded message to ask for confirmation that the receiving party wishes to receive the message. If it is confirmed, the MAP will play a recorded introduction and then, using a text to speech resource, will send the information content as a speech signal. "Message_start" and "message_stop" delineate the material which should be sent to the text to speech resource for conversion to a speech signal.

"Status" is another field which may be redundant in embodiments of the present invention. It is a field which the MAP completes for messages being returned to computing applications running services from non-MAP platform. In the primary embodiment of the present invention described herein, instead of obtaining dedicated status information in messages returned to the system from the MAP, the messaging system accesses the raw call records in the call record data store 2030. The MAP sends call record files here and enters status flags in the event of any change in status.

Referring to Figure 7, the call records process 2020 is as follows:

STEPS 3000, 3005: As described with reference to Figure 3.

STEP 7000: The user selects an HTML link identifying the call records process 2020.

STEP 7005: The CGI 2040 initiates the requested Call Records Process 2020 and returns a form to the user.

STEP 7010: The user returns a completed form.

STEP 3025: The process checks whether the completed form meets requirements and, if not, notifies the user and offers a limited number of retries (STEP 3035), as in the process of Figure 3.

STEP 7015: When the completed form is found to meet requirements, the process uses the information supplied to identify the user's call record data in the sorted

call records area of the call record database 2030. It then uses the data to send a page to the user showing a record of all the service instances they have generated, together with status information regarding whether a message has been successfully sent, has failed or is still being attempted.

In the normal course of operation, the MAP creates files of raw data on service instances it handles. These are each identified by the "service_instance_identify" field and stored in a raw call record data area (an FTP directory) of a database. It could be on the same computing platform 1000 as the Web server 1010, as shown in Figure 2, or it could be elsewhere. The messaging system of the present invention has access to the same directory. It sorts the raw call data according to user account number. It is these sorted records that the Call Records Process 2020 makes available to a user.

When the MAP is in the process of setting up a connection and transmitting a message to a recipient, it will output an update status message to the raw call records data directory, that is the FTP directory, whenever there is a change in status for a service instance. It uses the "service_instance_identify" to identify the relevant file. The Web server 1010 will similarly update files in the FTP directory when there has been a user input in respect of a service instance. The server 1010 also has access to this directory for the purpose of displaying call record data sorted by account number, in particular including status information, to the user. To avoid a backlog of updates occurring, the messaging system updates the sorted call records for display to a user by accessing the raw data whenever a logged in user selects a link and therefore moves between pages displayed at the Web browser 315.

Once the user is looking at a page displaying sorted call records, that page will not show new status data as it comes in from the MAP 150. For a user to check whether there has been a status update, it is necessary to ask for a refresh of the page, in known manner from the browser 315.

In a simple user interface, the information presented to the user to indicate connection status for all their messaging might be a table showing the name of each intended recipient, the title of each message, the date on which it was requested and a status flag having three conditions to show whether the message has been sent, has failed, or is still in the process of being set up. The status flag could be a coloured marker showing green for sent, red for failed, and amber for still in process of being set up, in the manner of traffic lights.

MAP 150

Referring to Figure 8, the MAP 150 sits between a data network 155, such as the Internet, and a telecommunications network 160 such as the PSTN. It comprises primarily one or more speech cards 110 and one or more digital line interface cards 105, managed by applications running on a PC 120. A network interface card (NIC) 130 provides connection to the data network 155, for instance by means of TCP/IP, and the DLICS 105 provide connection to a network switch or exchange 250, in this case a PABX, via an ISDN link 210.

In practice, the MAP 150 is capable of supporting many additional services and the following should be taken as a description of a platform capable of supporting embodiments of the present invention but also having functionality which may be unnecessary as far as embodiments of the present invention are concerned.

The speech cards 110 carry a set of algorithms for detecting, recognising, recording and generating speech signals and tones. The speech cards can also deal with modem signals. The PC 120 also carries some of the more sophisticated algorithms, such as the speech recognition and text to speech conversion algorithms. These are the resources which allow the MAP to recognise and convert between different signal media, such as between spoken messages, recorded voice messages, electronic mail messages and the like. Additionally, there are one or more FAX cards 165 for providing the capability of sending and receiving by facsimile.

The network interface card(s) 130 may of course support different communication protocols. As shown, the network interface card 130 supports TCP/IP and provides an interface to the Internet 155. This can give access to services provided by the MAP to any user having a connection to the Internet 155.

The digital line interface cards 105, the speech cards 110 and the FAX cards 165 are connected to each other by a (known) multi-vendor interface protocol link 100 and to one or more personal computers 120 by ISA links 115.

General control of operation of the MAP 150 is provided by a computer application running on the PC 120. In particular, this application controls flow of processing in the MAP 150. Actual process execution to provide services is primarily (but not exclusively) carried out by processors of the speech cards 110.

Referring to Figure 9, each speech card 110 carries seven speech processors. The first of these, a Motorola 68,000 speech processor 205 (referred to as "68K" in the Figures) , provides processing power and gives overall control of what is happening on the speech card 110. The other six speech processors are Motorola 56,000 speech processors 200 which provide DSP capability. The 56,000 processors 200 (referred to as "56K" in the Figures) provide some of the resources, particularly algorithms, used in running services on the MAP 150. These algorithms provide functions such as MF4 detection, playing out messages at 8 kbits/sec, recording messages at 64 or 8kbits/sec, and speech recognition.

Overall control of the MAP 150 is provided by the "MAP Manager". This is software which generally will reside on the PC 120. The MAP Manager has knowledge of the resources available, in terms of ISDN channels and the processors 200, 205, and it determines where resources will be made available to meet demand by applications. In particular, the MAP Manager has control over time switches 215 in the processors of the speech card 110 and allocates streams for processing by means of the MVIP 100.

Referring to Figure 10, the basic architecture of a MAP 150 can be multiplied up so that there are at least two PCs 120, connected to their respective MAP speech cards 110. The DLICs 105 may share a common platform, as shown. Each PC 120 is connected via the ISA bus 115 to a network interface card 130 and thus to the Internet 155 via a network link 125. Further processing power can be provided by a UNIX processor 300, also connected to the ISA bus 115 and hard disk storage 305 is provided. Further memory 310 is associated with the PCs 120. In this arrangement, one MAP 150 can communicate with, and exploit resources of, another MAP 150.

A MAP will now be described in more detail.

OUTLINE MAP SYSTEM DESCRIPTION

Referring to Figures 8 to 10, the Minor Applications Platform (MAP) 150 described below is a PC-based speech platform capable of providing multiple speech applications on multiple telephone lines. The heart of the system is the speechcard SC2 110 which provides speech and signal processing based on algorithms. The MAP 150 supports multiple digital links to public and private telephone exchanges using its Digital Line Interface Card (DLIC) 105. The MAP 150 supports resource sharing and is dynamically reconfigurable to run different applications on the same hardware. Internally the DLIC 105 and SC2 110 cards are interconnected using an industry standard speech bus 100 called Multi Vendor Interface Protocol (MVIP) and this capability may be used to build MAP systems using other commercially available cards. MAP systems may also be built in shelves with passive backplanes and use plug-in processor cards. Such backplanes may be split in which case MAP hardware and software resources may be shared between processors.

The MAP system has a single manager and can provide traffic statistics. Its operating system is UNIX. It has an archival facility which has direct links to DOS, WINDOWS and Netware.

THE MAP NETWORK SERVER

The MAP system 150 provides a 30 channel digital connection 210 to a public or private telephone switch 250, such as a PABX, using a variety of protocols. The MAP 150 allows connection to be made to a public exchange using signalling systems such as the DASS2 or A1B1 signalling systems, or to a private exchange using the DPNSS protocol. For the embodiment described below, the MAP 150 is connected to a private PABX using DPNSS which allows service provision based on TLI (and possibly SIC) codes. These codes have analogues in the PSTN where the known C7 signalling system is used, and thus services provided on the MAP as described herein can be implemented over the PSTN.

In this implementation 3 PCs 120, 300 are used to provide the MAP system. One PC 120A provides system management and connection to the PABX 250, whilst the other two PCs 120B, 300 provide advanced speech recognition and text to speech conversion services. A fourth PC (not shown) provides connectivity to the Internet 155 through an SMTP gateway, thereby enabling messages to be sent to and received from the LAN.

MAP Manager

The MAP manager provides control software for starting and terminating applications and co-ordinating their requests for services on line cards 105 and speechcards 110. The MAP manager resides on one of the PCs 120A and provides a screen-based user interface. From the manager screen the MAP system can be started, parked or shutdown and the state of resources examined. The manager uses a small number of tables which may be examined to control allocation of resources and application selection.

A MAP "Traffic Statistics" screen can be invoked from the manager and provides basic information on platform use by applications, notably number of calls/application, average call holding time and busy hour.

Referring to Figure 11, the MAP Manager 400 comprises control software 440 and a set of agents 405, 410, 415, 420, 425. Each agent is generally allocated to a functional entity, which may be hardware and/or software, and provides a control link thereto. Hence there is a DLIC agent 405, a speech card agent 410, a fax agent 415, a text to speech agent 420 and a speech technology applications processor agent 425. The DLIC, speech card and fax agents 405, 410, 415 are provided with Applications Programming Interfaces (APIs) 430 and with plural drivers 435 for driving the relevant pieces of hardware (cards). Each API is provided by the suppliers of the relevant hardware cards. The MAP Manager control software 440 provides overall control of the agents and acts as a sophisticated message switch in that substantially all messages go via the MAP Manager control software 440.

The drivers 435 receive incoming data and control outgoing communications for example by controlling interrupts on the relevant bus.

The MAP Manager 400 provides overall resource allocation. To do this, at start up, it constructs a world view. This is checked and updated at intervals. It obtains the world view through its resource agents. For instance, referring to Figure 12 which shows a representation of a small part of the MAP Manager's world view, the speech card agent 410 may have reported that two digital signal processors 500, 505 are loaded respectively with a multi-frequency detection algorithm 510 and an 8K coding scheme 515. These two digital speech processors 500, 505 will therefore deal with signals on the MVIP 100 which are incoming and outgoing respectively.

MAP Processor

MAP can be potentially hosted on most standard PCs with a 33 Mhz 486 PC or better and having a minimum of 8 Mbyte of memory. The standard processor used within the BT network is the Compaq 5133, a Pentium based machine having 5 full length ISA slots for MAP and other hardware cards.

As an alternative, a processor card manufactured by Aculab plc can be used. This is intended for use in a shelf unit where the need is to provide a multi-processor MAP system in a single enclosure with minimal non-MAP components. Outline features include

- 80486 DX CPU operating at 25,33 or 50 Mhz
- Internally clock doubled parts may be fitted
- Pentium ready
- 256 kByte cache
- ISA Bus
- On-card Ethernet LAN interface
- Serial and Parallel Ports
- IDE Hard Disk and Floppy Disk Interface
- Keyboard Interface

MAP Algorithms

Speech algorithms for use on the speech cards 110 and processors 120, 300 include:

64kbit/s speech I/O
8 kbit/s speech I/O
Silence detection
DTMF Detection
DTMF Generation
Network Tone Detection
ASIIWR CDHMM Isolated Word Recognition Yes/No, Digits, alphas
ASIIWR Rapid Vocabulary Generation
STAP CDHMM Connected Recognition
STAP Rapid Vocabulary Generation
Fast/Slow Speech
Pulse Recognition

Noise Interrupt
Text to Speech
V21/V23 Modem
Audioconferencing

MAP Speech Card

The Speechcard is manufactured by Aculab plc and provides a UNIX-based environment for algorithm execution, with a VXWorks real-time kernel. It contains six Motorola 56001A DSPs clocked at 27MHz for digital signalling processing each of which connects to the MVIP bus. The on-board 40 MHz Motorola 68000 coordinates communication between the DSPs and the MAP host processor over the ISA bus. It is also used in some instances for further algorithmic processing.

The card is EMC approved and also provides a single analogue port for local speech recording/playback facilities.

MAP Operating System

The MAP operating system is System 5 release 4.4 UNIX. This is commercially available as Consensys, or Unixware V1.0. The MAP system could of course be ported to other systems such as Unixware V2.0 or LINUX.

The MAP system must also interwork with the speechcard host operating system, VxWORKS, a UNIX system with a real-time kernel. One copy is required per speechcard.

MAP Shelf Unit

This is manufactured by Aculab plc and contains an 18 slot passive ISA backplane organised into four processor sections as 5 + 5 + 5 + 3. Processor sections may be

interconnected to make larger processing units. It is primarily used to contain multi-processor MAP systems and archival systems and contains

- Mains Power Supply
- Good Ventilation with fan failure monitors
- Internal provision for 4 Winchester Disks
- Floppy Disk Access
- Twin DAT Drive Access

MAP Digital Line Interface Card

Primary Rate ISDN connection is provided using Aculab's "E1" card. There are 30 and 60 channel versions available on an ISA PC bus. This E1 card supports different types of speech bus 100, including MVIP. It is EMC approved and is widely used and approved for host-independent connection in many countries worldwide.

In the MAP system, as discussed above with reference to Figure 4, DLICs 105 are controlled by a DLIC Agent 405. Under the control of the Manager 440, speech from a timeslot on an ISDN link 210 is switched by the DLIC 105 onto an MVIP timeslot and stream and, using a similar switch on the SC2 110, the manager directs the speech to a given DSP 200.

Referring to Figure 10, DLICs 105 in MAP systems 150 can also be "daisy chained", allowing one MAP system to pass calls to another, only the first MAP system 150A being connected to the host exchange. This avoids a situation as shown in Figure 9 where, in order to get sufficient MAP service capacity, a customer is supporting direct ISDN connections to three different MAPs 150. This might otherwise be the result where a customer has an ISDN connection for 30 channels on one link but one MAP system 150 cannot provide all the required resource. To "daisy chain" MAPs 150, it is simply necessary to substitute a sixty channel DLIC 105 for the thirty channel DLIC 105 used normally in a MAP, in the first and second MAPs 150A, 150B, and to use the second and third MAPs 150B,

150C for overflow connections via the DLIC agents 405 of the first and second MAPs 150A, 150B.

MAP SERVICE PROVISION

Referring to Figures 10 and 14, requests for services using the MAP may be received either on one of the ISDN links 210 or, via a network interface card 130 and an SMTP gateway process 725, from a computer-based link such as a UNIX socket connection, a LAN, an Intranet or the Internet 155.

Although in the embodiment of the present invention described above, the computer platform 1000 outputs an SMTP message to the MAP, it could equally use a different format and could be delivered to the MAP 150 for instance by means of an ISDN connection. Hence the following description includes the functionality of the MAP in the scenario that a service request message is received by ISDN connection.

To respond to a request for service, the MAP 150 will need to run a computer-based application 730,735 to manage the service, and it will need to allocate the resources necessary to run the service. These resources may include at least one ISDN time slot 720. Service requests coming in on an ISDN link 210 will necessarily already have an ISDN time slot 720 and the MAP allocates these directly to the relevant applications 730, 735. Service requests coming in from the computer-based link on the other hand are placed in a Generic Service Queue 710 and the various applications 730, 735 which are relevant to these services read out their own service requests from the Generic Service Queue 710.

It will sometimes be the case that an application 730, 735 will need to call on other applications to fulfil the service. In these cases, the application which is running will effectively issue a service request 715 which goes to the Generic Service Queue 710 until the newly requested application reads it out again. This might occur for instance where a customer has designated that a failed paging attempt should be replaced by fax notification.

Applications 730,735 on MAP speech channels 720 can in general be caused to run in two ways.

1. Immediately on system startup. This mode is used, for example, as the basis for provision of speech services requested via a computer-based link. If a speech service request is received via a computer-based link, there is no speech channel 720 directly involved. One needs to be provided as quickly as possible. A minimum of channels 720 are therefore permanently dedicated to each service, though an application can acquire further lines from a free pool. Such applications 730,735 are normally outgoing only e.g. notification and voicemail delivery services, or booked services such as conferencing.

2. On receipt of incoming call. In this case, there is necessarily a speech channel 720 associated with the service request. The MAP system allows the application to be selected by channel of entry or DDI number, both in conjunction with the SIC. In the present embodiment, the DDI method has been adopted as this allows line plant to be used according to demand with consequent higher line utilisation. This facility would be used for services such as telephone access to e-mail.

As mentioned above, more sophisticated services may be conveniently built up on more basic services - for example, should it be desirable to send a FAX as part of one service, the document for transmission can be sent to a FAX service for transmission to the intended recipient. Such delivery services are non-real time and by their nature do not call upon other services. Requests for such services are queued, with the relevant MAP applications taking requests from the generic service queue 710.

The Generic Service Queue 710 can be divided into application types so that there are multiple service queues 710. A number of service channels 710 can be pre-

allocated to each service queue 710 and the number of service channels for a particular service queue can be determined from observed traffic levels.

An Internet address can be associated with each generic service queue. Requests can then be added to the queue by sending messages to this address. These messages need to contain all the information required for the MAP application to subsequently enact the request, and are thus best computer generated. Client software to originate such messages is described above, and is the basis on which 'point and click' telephony and messaging services can be provided. However, it should be noted the MAP applications can themselves request such backbone services. For example, if a user of the message collection service wants to receive a copy of an e-mail message by FAX this can be achieved by the voicemail application making a FAX request. This is the basis of the collection services.

MAP Resource Allocation

This is further described below. However, resource allocation in general is as follows:

- *Line Allocation* is specified in configuration files which may be dynamically updated while the system is running. Lines may be classed as reserved for outgoing (free pool), fixed for incoming, or available for either. Lines may be allocated either statically or dynamically.
- *Application Allocation* is either by point of entry (ordered) or by DDI digits for service requests incoming by ISDN channel. For incoming SMTP messages, it is simply done by interpreting the service requested and directing the request to an "IMMEDIATE" application. Applications are normally run by the manager but, like hardware resources, may be distributed between MAP systems sharing a common MVIP.
- *Algorithm Allocation* may be either static, dynamic or a combination of both. There can be only one algorithm per DSP, but (depending on the algorithm) one DSP may be able to serve more than one channel. There is no fixed

algorithm/DSP association and this will typically change according to the demands of applications.

- With *static* resource allocation, all resources are allocated at the time of call arrival and will be always available even if unused. With the *dynamic* scheme, resources are only allocated when needed which enables more channels to be supported for a given level of DSP resource - there is however a probability of demanded resources being unavailable which must be suitably catered for within the controlling speech application. This situation can be handled in the MAP by the application requesting resources in advance and manually releasing them when no longer required.

INCOMING PSTN CALL PROCESS AT MAP 150

Referring to Figures 10 and 11, a user initiates a dialled connection to the MAP 150, using a telephone 320 connected to the PSTN 160. The call will be received at a DLIC 105 and a driver 435 for the relevant DLIC agent 405 will receive an event on a timeslot "X". The DLIC agent 405 has its own processing capability and will instruct the driver to acknowledge the event on timeslot "X". The DLIC agent 405 will have also instructed the relevant DLIC 105 as to what protocol is being used. The DLIC agent 405 then tells the MAP Manager control software 440 that there was an incoming event.

The MAP Manager control software 440 recognises that it will have to raise a computer-based application to provide the requested service. To do that it needs to obtain the DDI digits which will identify the service requested by the user. The MAP Manager 400 will therefore tell the DLIC agent 405 that it needs the DDI digits.

Once the MAP Manager 400 has the DDI digits, it will access an interpretation table stored on disc 305. In practice, the disk 305 will store any frequently accessed data in local memory, for read access only. (Nowadays this is a standard feature of many disk storage systems. It is controlled by a disc controller provided by the disc supplier.) The DDI digits, once interpreted by use of the table, will

identify to the MAP Manager the service required. The MAP Manager 400 needs to translate this into the application(s) needed to provide the service, and thus into resources, including algorithms, required. It therefore looks up another interpretation table, in the same manner as before, in order to find out what resources and algorithms are going to be needed to run the applications necessary to provide the service. (In practice, one interpretation table may provide both service and resource data.)

Knowing the resources, including algorithms, it will need, the MAP Manager will now look at its own world view, constructed on start-up, to find a functional and available digital signal processor 200 and spare ISDN channels 210, with which to run the relevant application. Having found the relevant resources, the MAP Manager 400 books the resources by entering a flag in the internal table which is the world view and sends messages back to the relevant agents 405, 410 to allocate switch settings 215 and to launch the relevant application.

Referring to Figure 13, the first task of the application is usually to answer the incoming call. The application 600, using the relevant API 605, now tells the MAP Manager that the call should be answered. The MAP Manager 400 tells the DLIC agent 405 to answer the call and the DLIC agent 405 passes the message on to the application via the MAP Manager 400. The MAP Manager 400 now reports back to the application 600 that the call has been answered or that there has been a failure. At this point, the user who has initiated the call hears the ring tone leave the line.

The application next instructs the MAP Manager 400 to play a file. The MAP Manager 400 does this by means of the speech card 110 and by passing a request to the SC2 agent 435. Once notified that the file has been played, the MAP Manager 400 sends notification back to the application 600.

The above describes the first steps in the running of an application 600 to provide a service using the MAP 150. The application 600 will continue to control provision of the service until it has concluded. At this point, the MAP Manager

400 will send an acknowledgement back to the DLIC agent 405 and the DLIC agent 405 will clear the line. This means the switches 215 will be opened and the resources will be "torn down", for instance by marking the digital speech processors 200 as available again in the internal table of the MAP Manager 400.

It will be recognised that the DLIC 105 carries relatively little data. There is significantly more data on the speech cards 110. The 68,000 processor on a speech card 205 runs the VXWorks UNIX kernel 170 in real time, by means of which it can run VX Works. If the speech card 110 receives a play request from the MAP Manager 400, it is VX Works which handles it. That is, it runs a file off disk without involving the main processor. (This is a known use of UNIX.)

It is possible to do the messaging between the MAP Manager 400 and its agents 405, 410, 415, 420, 425 via common memory or via a network interface card 130. VX Works will run as though it were stand-alone, using a network interface card 130.

It will be recognised that although the speech cards 110 do a significant amount of processing, processing can also be distributed elsewhere. A constraint is that although applications can sit in many different places, all speech signals should be carried via the MVIP bus 100. Consequently, speech databases should sit next to their speech cards 110 to avoid carrying speech signals on the Internet, or other relatively unsuitable highway.

INCOMING SMTP MESSAGE PROCESS AT MAP 150

This process is the basis for the embodiment of the messaging service of the present invention described above particularly with reference to Figures 3 to 7 and further described under the heading "E-Mail Delivery Service" below.

Referring to Figure 14, service request messages enter the system through the SMTP gateway 725 from the computing platform 1000. The service request is physically transferred to the system and (usually) placed on the relevant generic

service queue 705, 710. This gateway 725 relates SMTP addresses to generic service queues 705, 710. The addresses of the available generic service queues 705, 710 are of the form

<user>@<service>.map.cartoon.bt.co.uk

where <service> is the name of the generic service queue. It should be noted that the <user> field may not be used or may have different interpretations in different services. Messages to non-existent services are discarded. Messages to services which require immediate delivery will be sent directly to the relevant service channel queue 730, 735.

So far as launching applications and delivering resources for services is concerned, the process is the same where a service request has been received at the SMTP gateway 725 as described above for the "Incoming PSTN Call Process". That is, the MAP manager 440 will use its various agents as described above with reference to Figure 11.

PROBLEM SCENARIOS AT MAP 150

User Rings Off

The above describes service provision which runs to its normal end. It may be that the user rings off while a file is still being played. In this event, the DLIC agent 405 will notify the MAP Manager 400 which in turn notifies the speech card agent 410 that the file should be stopped. The speech card agent 410 will tell the relevant speech card 110 to stop the file and report back to the MAP Manager 400 which in turn sends an acknowledgement back to the DLIC agent 405, and notifies the application that the line has cleared, once any house-keeping activites are completed, such as billing. Once complete, the application terminates which results in tearing down the resources, as before.

Wrongly Distributed Resource

There are other problems which may arise. For instance the MAP Manager 400 may find that it does not have sufficient resource, such as "play". In these circumstances, the MAP Manager 400 will seek to change the available resources and therefore to change its world view. It may do this by deleting "recognition" from a processor, as long as it is not in use, and substituting "play" resource.

Resource Failure

In a different problem scenario, there may be functional difficulties, for instance with a digital signal processor 200. Say a speech card 110 was providing a play file. This is done in discreet amounts. The digital signal processor 200 should respond in a given time (to VX Works) to say it has played a required section of a file. If a digital signal processor 200 has "died", VX Works will never receive an answer. In these circumstances, the speech card agent 410 will reload the code to the digital signal processor 200. It will, in practice, try this reload three times. If the speech card agent 410 has not received any response from the digital signal processor 200 after three reloads, it will report the digital signal processor 200 to the MAP Manager 400 as being of doubtful integrity. The MAP Manager 400 will select another DSP 110 and load the relevant algorithm. The MAP Manager 400 will subsequently only use the digital signal processor 200 marked as of doubtful integrity in emergencies.

If then another digital signal processor 200 fails, the MAP Manager 400 has a potentially serious situation. It now issues a message to the control console (accessible from the MAP Manager screen), and issues a serious fault indicator on the serial line for that purpose. It may not be the digital signal processors 200 which are failing but VX Works for instance. The MAP Manager 400 will also take the relevant card or cards out of service and use spare resource. At this point the MAP Manager 400 reboots the suspect card. This forces the card to go through a "power-on self test". The card will then be proven either serviceable, in which case the MAP Manager will put it back into use, or failed.

Certain applications require large amounts of resource. For instance, speech recognition would create a bottleneck for the digital signal processors 200, 205. Referring to Figure 11, the text to speech (TTS) agent 420 accesses resources in a dedicated processor via a network interface card 130. The TTS agent 420 will send a sentence to the dedicated processor. It receives back a 64K speech file which the TTS agent 420 runs and plays through a digital signal processor 200.

Running the digital line interface card 105 is relatively easy in view of the lower data quantities. More of the management effort tends to lie with the speech cards 110. The processor on the MAP host, running UNIX, is usually underutilised. This processor can be used to host text to speech or STAP algorithms. It should be noted, therefore, that the TTS agent 420 is not controlling hardware in the manner of the other agents 405, 410, 415, 425 but is controlling processor resource. The TTS agent 420 therefore has to allocate resource in a different way. The speech card agent 410 measures resource in terms of channels on the digital signal processors 200 and how much 68K processors hire an application will take. For instance, for multi-frequency tone recognition, ten channels would need to be allocated and 2% of a UNIX controlling process. Looking at the TTS agent 420, it is necessary to look at processing resource and it will be measured in megabytes and channels, for instance 32 megabytes and seven channels. The TTS agent 420 is directly connected to the speech card 110 by socket, this providing a real time connection. The MAP Manager will therefore instruct the speech card agent 410 that it needs 64K play capacity and a socket interface. TCP/IP runs on the Ethernet at about 2 megabits per sec which is the equivalent of about 27 channels. Hence the socket arrangement avoids a bottleneck at the PC bus point which would only be able to provide about 12-18 channels off disk.

The STAP agent 45 provides the inverse process to the TTS agent 420. Speech is received through the speech card agent 410 which does some pre-processing. It then sends speech to the STAP agent 425. The STAP agent 425 returns a recognition result and the speech card agent 410 notifies the MAP Manager 400 accordingly.

OUTGOING FILE PROCESS FROM MAP TO MESSAGING SYSTEM

Some backbone services potentially require an element of communication with the requesting user. This might be to provide a booking reference number, to advise of problems in being able to contact some parties, or perhaps to advise of an inbound communication. To support such services, the MAP SMTP gateway 725 must be able to create/forward outbound messages from the MAP platform 150 and some client software is also desirable to receive and act upon such messages. This client software would normally be iconised until activated by message arrival when it should interrupt the user to request what action should then be taken. This requires extensions to the MAP SMTP gateway 725 and client software to act on such messages. (Client software is already available which can be easily modified to accept such messages.)

However, for the purposes of the embodiment of the present invention described above with reference to Figure 7, the MAP primarily needs to output files via an FTP link to the call records database 2030 for use by the call records process 2020 associated with the user's Web server 1010.

MAP BACKBONE SERVICES

These services are important in enabling more sophisticated services to be provided, incorporating within them one or more backbone services. These can be addressed by client software. (Other "backbone" services may also be provided to enable facilities such as user registration and call accounting). For the purpose of the present specification only one of these backbone services is described below. This is the E-Mail Delivery Service.

E-Mail Delivery Service

The queue for this service takes as its input a service request message (ie in text) delivered from the user's computing platform 1000. It attempts to deliver this

message by speech to one or more specified telephone numbers. The message may be delivered by the voice determined by the service address chosen according to the following table:

Service Address	Voice
<user>@ttsm.map.cartoon.bt.co.uk	British Male
<user>@ttsf.map.cartoon.bt.co.uk	British Female

The service comes in two forms. If the user field is present in the address and is all digit, it is interpreted as the telephone number which should be dialled from the United Kingdom (UK) to contact the (single) receiving party, with the body of the message being taken as the unstructured whole of the communication. If the user field is absent, the message body is taken as being structured potentially to enable distribution to many parties.

The MAP 150 can make a number of attempts at delivery of a message to a given party and, if the message cannot be delivered, a notification message can be dispatched to the originator giving the reason for non-delivery. If the target telephone line is answered, there is a simple service announcement followed by a query as to whether the message can be delivered. Once delivered, the customer can have the option to hear the message again or to receive a FAX copy. If delivery was turned down, a simple message can be sent to the originator detailing the non-delivery.

In the embodiment of the present invention described above with reference to Figures 1 to 7, use of the "E-mail Delivery Service" is as follows:

The MAP platform receives the request message and processes it, using the text of the message to synthesise a voice message. It then attempts to call each of the telephone numbers it has been passed in the request message.

If the call is answered in a predetermined time, the receiver is requested to press a DTMF (Dual Tone Multi Frequency) button and the synthesised message is

delivered. If no DTMF button is pressed, or the call was not answered in the predetermined time, no message is delivered and the MAP will attempt to redeliver it a number of times, after specified intervals.

The status of all these messages can be recorded by the MAP platform 400 and used as the basis for charging the customer. As mentioned above, records at the WWW server can also be updated from the MAP platform 400. These records can be used to build WWW pages for use when a user requests information about the state of her request.

It is not necessary to use a WWW browser. It could be replaced with with a specialist client for instance, which may offer the user greater versatility. Similarly, the WWW server could be replaced with a specialist server. The client server could make use of portable code such as JAVA applets or Microsoft ActiveX.

The client could be replaced with an automated device such as a monitor. this would then send the request to the server when a monitored device reached a specified state.

The MAP platform 400 could be replaced by other software/hardware capable of text to speech conversion and call set up. The PSTN may also be replaced by another form of network which can carry voice, including the Internet.

The MAP platform 400 could be replicated and located at goegraphically separate locations. Request messages can then be routed to an appropriate instance of the MAP platform (or its equivalent) according to a routing algorithm located at the WWW server or at another location, for instance a computer located between the WWW server and the MAP platform 400. This would enable network loading, cost and value to be adjusted to the benefit of the service providers and the customer.

MAP applications can generate log files which describe actions taken by the platform for a particular call. Additional information may be held in these log files

to provide information for billing purposes. For example, the 'servicename' and user account number can be included; this information could be included as part of the structured messages so that layered services such as 'PhonE-mail' which call upon other, more basic services can pass the essential billing information forward in a way which is compatible with GUI originated messages. The log files once created would normally be dispatched to a specified mail address where they would be analysed to extract information for billing purposes. If this is to be handled on-platform the receiving process may take the form of a service queue. Potentially, billing information can be made available to the user via WWW or other network connection.

CLAIMS

1. A messaging system for use in providing a communications messaging service, which system comprises:

i) an input for receiving a messaging service request, the service request comprising destination information and message content information;

ii) a user data interface and store for receiving and storing user input data;

iii) request processing means for processing a received service request; and

iv) an output for outputting a processed service request in a format for use by communications connection set-up means in transmitting the message content information, or information identified thereby, as an output message to at least one destination identified by the destination information.

2. A system according to Claim 1 wherein the destination information in a messaging service request may comprise information identifying more than one destination, the communications connection set-up means being adapted to respond to a processed service request comprising destination information identifying more than one destination by sending common message content information to each of the destinations identified.

3. A system according to Claim 2 wherein a messaging service request is structured to include at least two data fields and the request processing means is adapted to process the request as a request comprising information identifying more than one destination in the event that a specified one of the data fields is empty.

4. A system according to any one of the preceding Claims which further comprises means for inputting precomposed message content to the user data store and the request processing means is adapted to refer to the user data store

and to incorporate precomposed message content in a processed service request in the event that a received service request comprises an identifier for the precomposed message content.

5. A system according to any one of the preceding Claims wherein the input for receiving a message service request is adapted to receive message service requests having a format which complies with an electronic mail protocol.

6. A system according to any one of the preceding Claims, which system is at least partially provided by computing platform providing a screen-based interface for a user to input message service requests.

7. A system according to any one of the preceding Claims, further comprising the communications connection set-up means, said communications connection set-up means comprising means to generate and store transmission status data in respect of output messages it transmits, wherein the system further comprises means to select and access one or more items of transmission status data so stored such that said item(s) can be made available to a user.

8. A system according to Claim 7 wherein the transmission status data is updatable during the operation of the communications connection set-up means to set up a connection for the purpose of transmitting an output message, such that a user can obtain information concerning the status of the connection being set up at a time after making a messaging service request and prior to an output message being transmitted.

9. A system according to any one of the preceding Claims, which system further comprises registration means for use by a user to enter customer-related data, which registration means has means to allocate a customer identifier in response to input by a user of preselected data instances, the user data store being adapted to store data sorted in accordance with customer identifiers.

10. A system according to Claim 9, which system is provided with a user login means, means to relate a customer identifier to a user on login, and means to select transmission status data for downloading at least in part according to the customer identifier relevant to a logged in user.

11. A system according to any one of the preceding Claims which further comprises means for incorporating scheduling data in a messaging service request for use by the communications connection set-up means in scheduling transmission of output messages.

12. A system according to any one of the preceding Claims, which further comprises means for inputting directory data to the user data store for customer-related data, the directory data comprising a set of user destination inputs and a set of destination identifiers for use by the communications connection set-up means in transmitting an output message, wherein the request processing means includes means for accessing the directory data and means to substitute a plurality of destination identifiers in place of a single user destination input in a received service request.

13. A system according to any one of the preceding claims wherein the processed service request which is output by the system for use by communications connection set-up means has a format which complies with an electronic mail protocol.

14. A system according to any one of the preceding Claims, which system comprises software for controlling computing platform which provides client/server capability in a client/server architecture.

15. A system according to any one of the preceding Claims, which system comprises software for controlling computing platform which is connected to a communications network adapted to carry communications according to the Transmission Control Protocal and the Internet Protocol.

16. A system according to Claim 15 wherein the computing platform is connected to the communications connection set-up means by means of a real-time data connection.

17. A system according to any one of the preceding Claims, wherein the communications connection set-up means is provided with text to speech translation capability and is connected to a telecommunications network such that the communications connection set-up means can translate message content information in a received service request from a text format to a speech format and transmit an output message to the telecommunications network.

18. A system according to any one of the preceding claims , further comprising the connection set-up means, said connection set-up means comprising:

i) a voice network interface, for providing access to a network for carrying voice signals;

ii) a data network interface, for receiving processed service requests;

iii) a resource interface for access to resources for use in providing communications services, including at least one speech-related resource from the group comprising voice recognition, recordal of incoming sound signals and transmission of outgoing sound signals;

iv) interpretation means for use in relating a processed service request to a computer-based application for use in provision of that service; and

v) initiating means for initiating the running of a computer-based application related to a processed service request

so as to generate a common voice message output to more than one location in the network carrying voice signals.

19. A system according to any one of the preceding claims wherein the network for carrying voice signals comprises a telecommunications network, providing voice channels in which one or more connections are established and maintained to support a communications session throughout its duration.

20. A system according to Claim 19 wherein each of said connections is established over a fixed route for the duration of a session.

21. A system according to any one of the preceding claims which further comprises queuing means for queuing received service requests, said initiating means responding to received service requests in an order determined by the queuing means.

22. A method of transmitting a message on a communications network, which method comprises the steps of:

i) receiving a structured service request; and

ii) processing the service request to generate a processed service request which is suitable for use by a communications connection set-up means in transmitting a message,

wherein said processing includes reference to customer-related data for use in the processing.

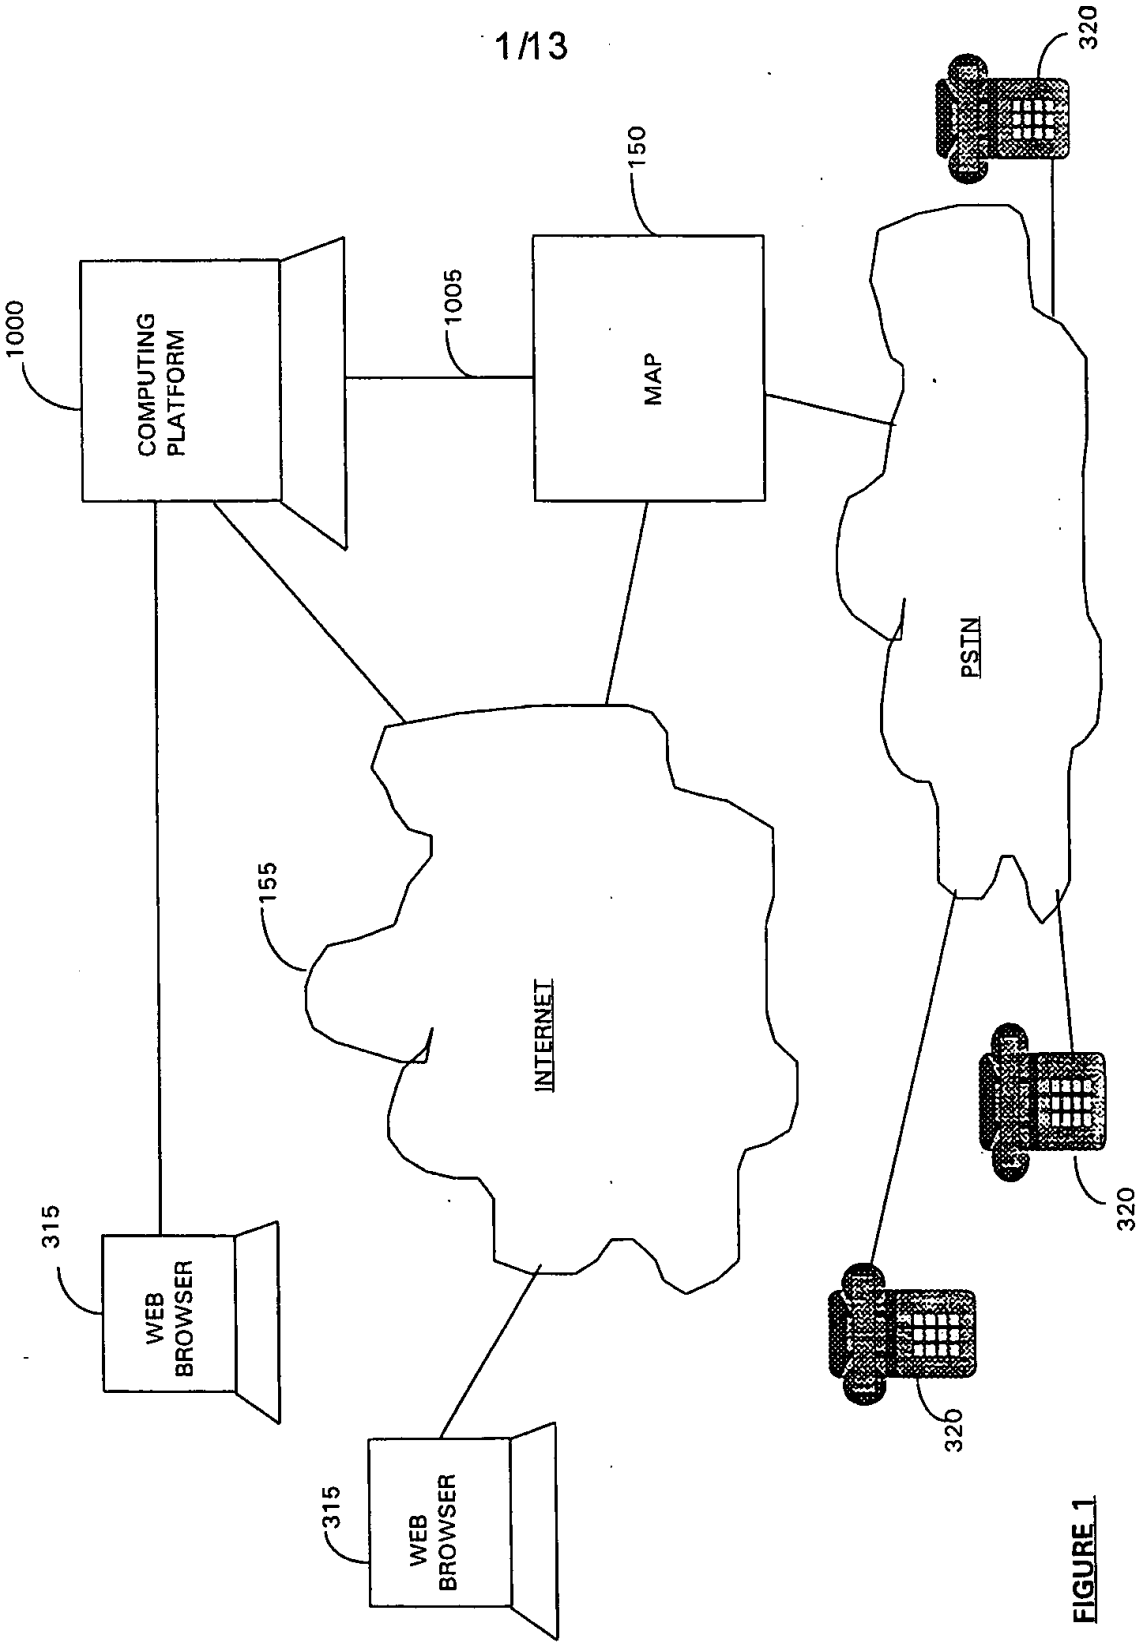
1000
COMPUTING PLATFORM
1005
150
MAP
320
PSTN
155
INTERNET
315
WEB BROWSER
315
WEB BROWSER
320
320

FIGURE 1

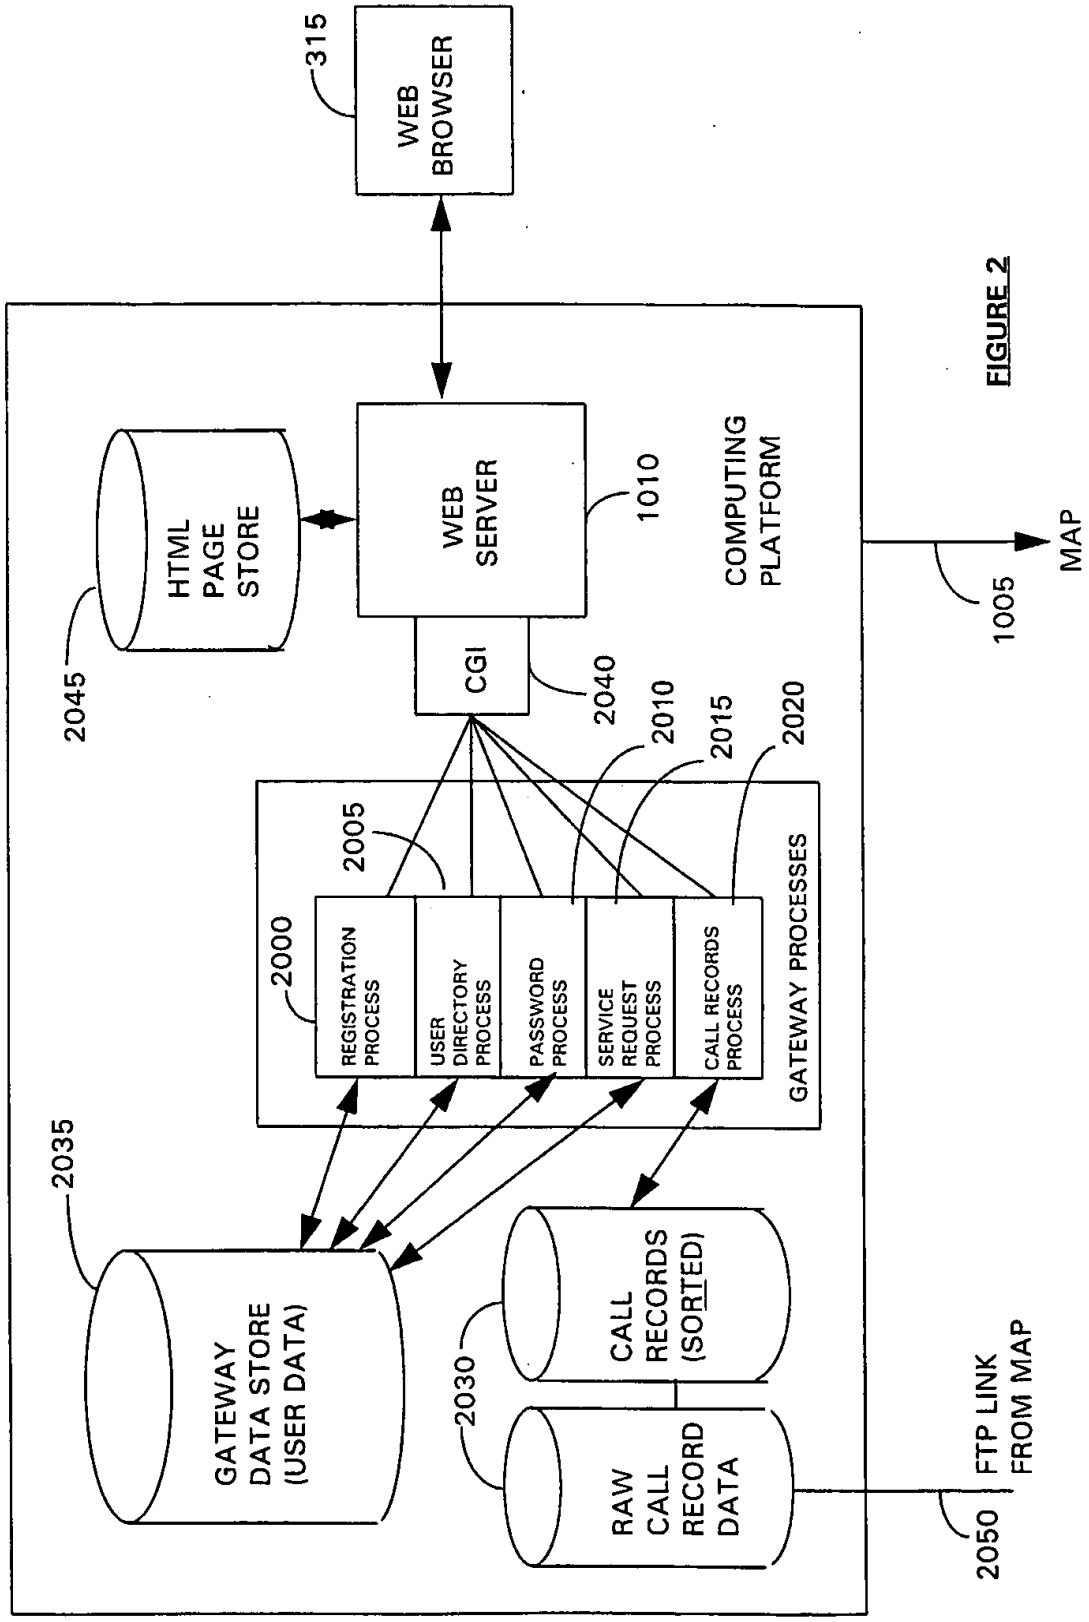
315
WEB BROWSER
WEB SERVER
1010
HTML PAGE STORE
2045
CGI
2040
COMPUTING PLATFORM
1005
MAP
2000
2005
2010
2015
2020
REGISTRATION PROCESS
USER DIRECTORY PROCESS
PASSWORD PROCESS
SERVICE REQUEST PROCESS
CALL RECORDS PROCESS
GATEWAY PROCESSES
2035
GATEWAY DATA STORE (USER DATA)
CALL RECORDS (SORTED)
2030
RAW CALL RECORD DATA
2050
FTP LINK FROM MAP

FIGURE 2

FIGURE 3

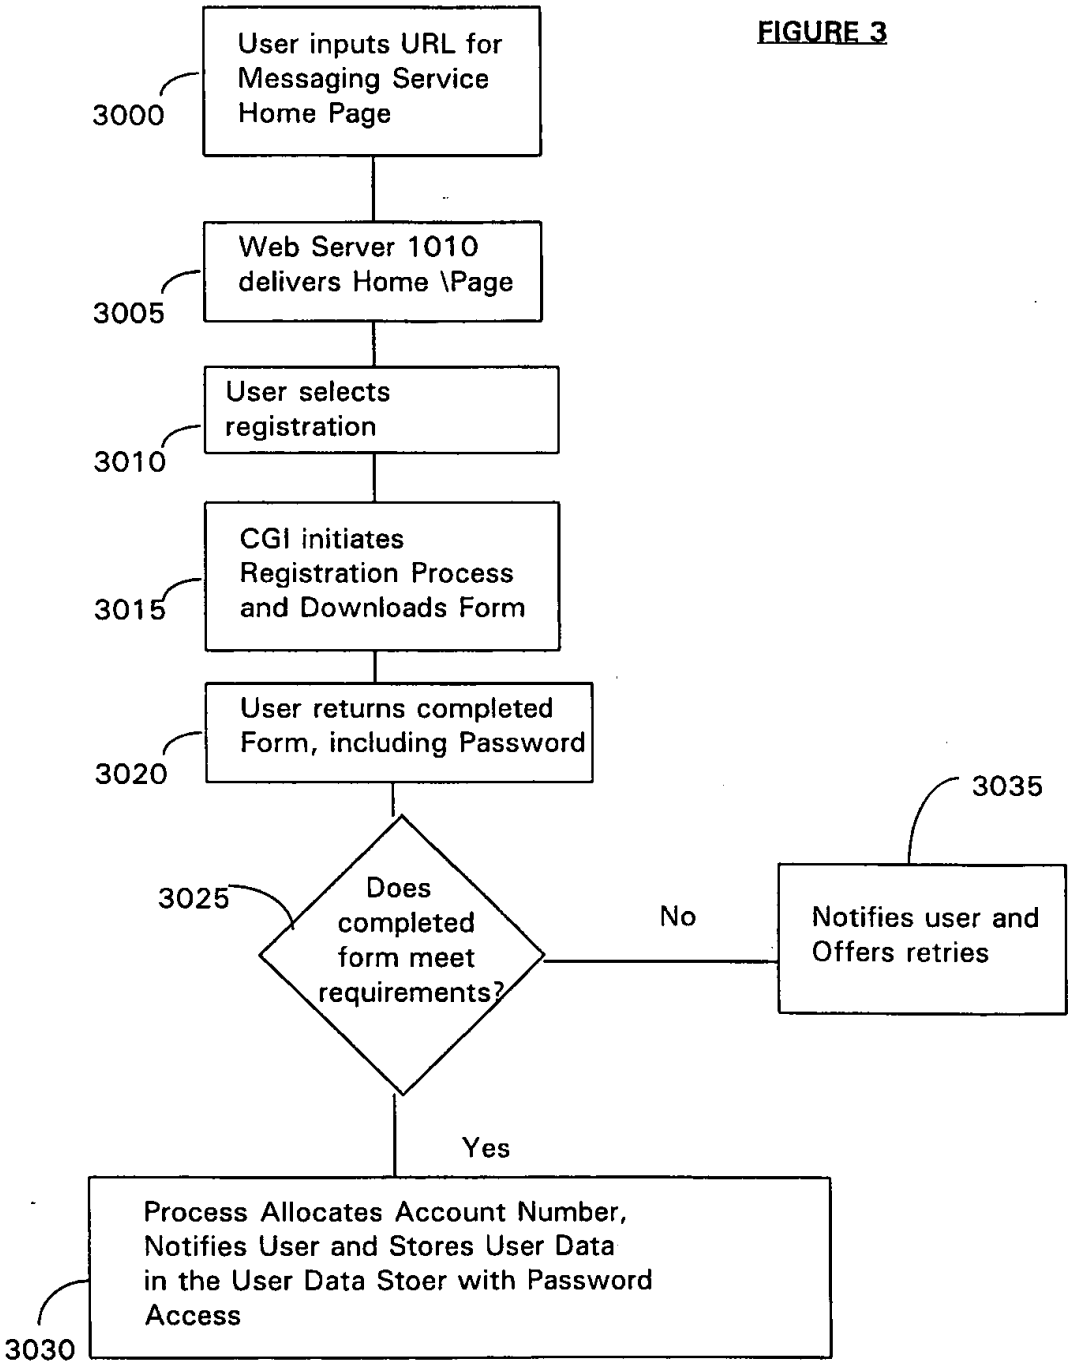
User inputs URL for Messaging Service Home Page
3000
Web Server 1010 delivers Home \Page
3005
User selects registration
3010
CGI initiates Registration Process and Downloads Form
3015
User returns completed Form, including Password
3020
Does completed form meet requirements?
3025
No
Notifies user and Offers retries
3035
Yes
Process Allocates Account Number, Notifies User and Stores User Data in the User Data Stoer with Password Access
3030

FIGURE 4

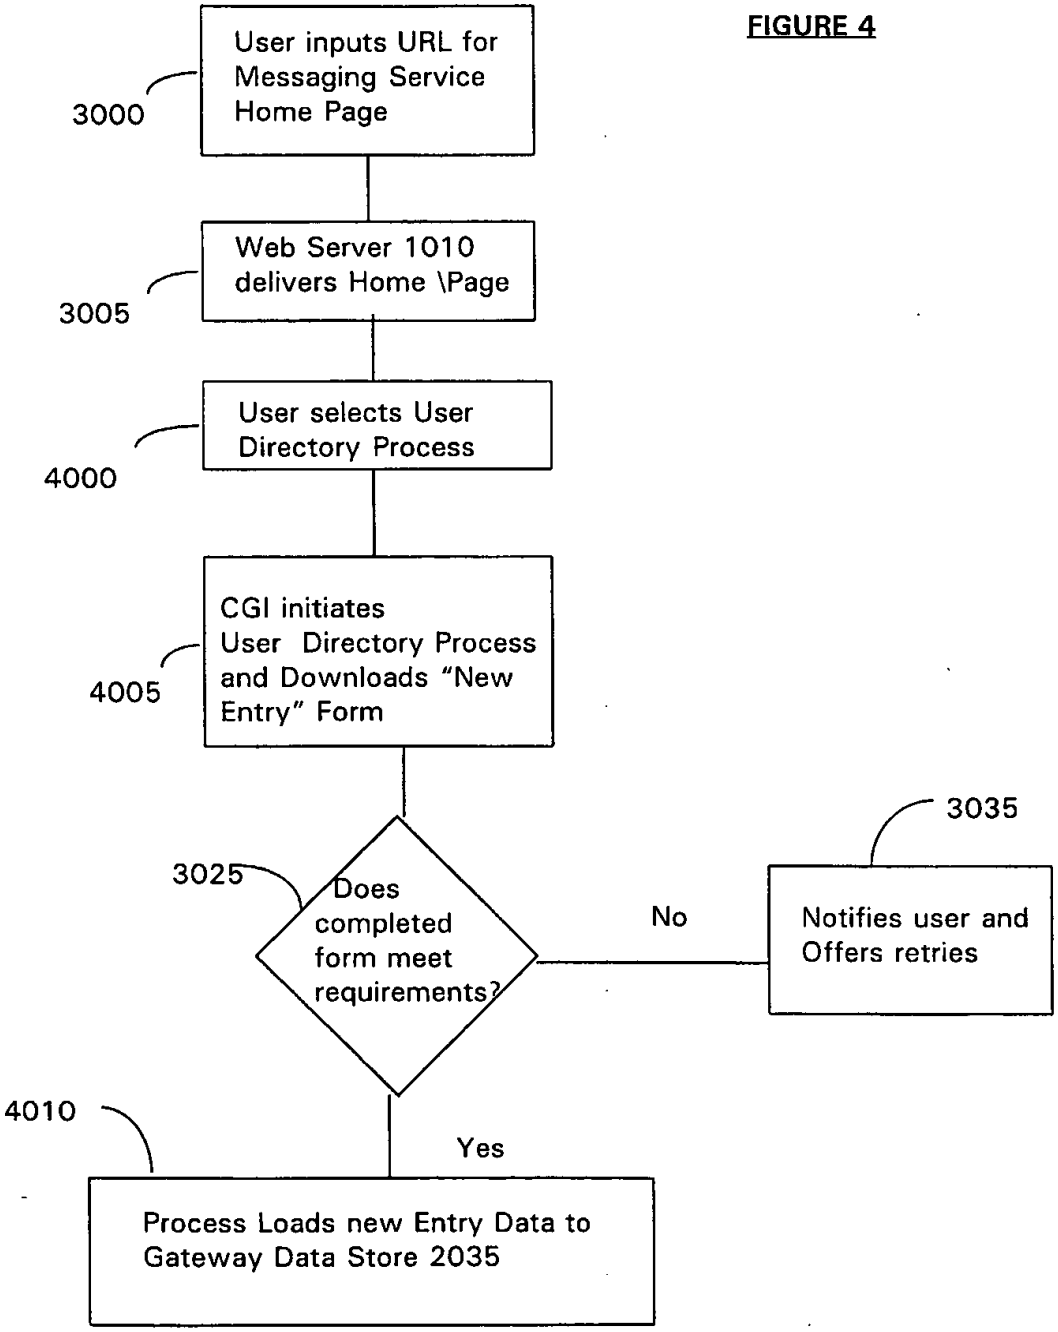
User inputs URL for Messaging Service Home Page
3000
Web Server 1010 delivers Home \Page
3005
User selects User Directory Process
4000
CGI initiates User Directory Process and Downloads "New Entry" Form
4005
3025
Does completed form meet requirements?
No
3035
Notifies user and Offers retries
Yes
4010
Process Loads new Entry Data to Gateway Data Store 2035

FIGURE 5

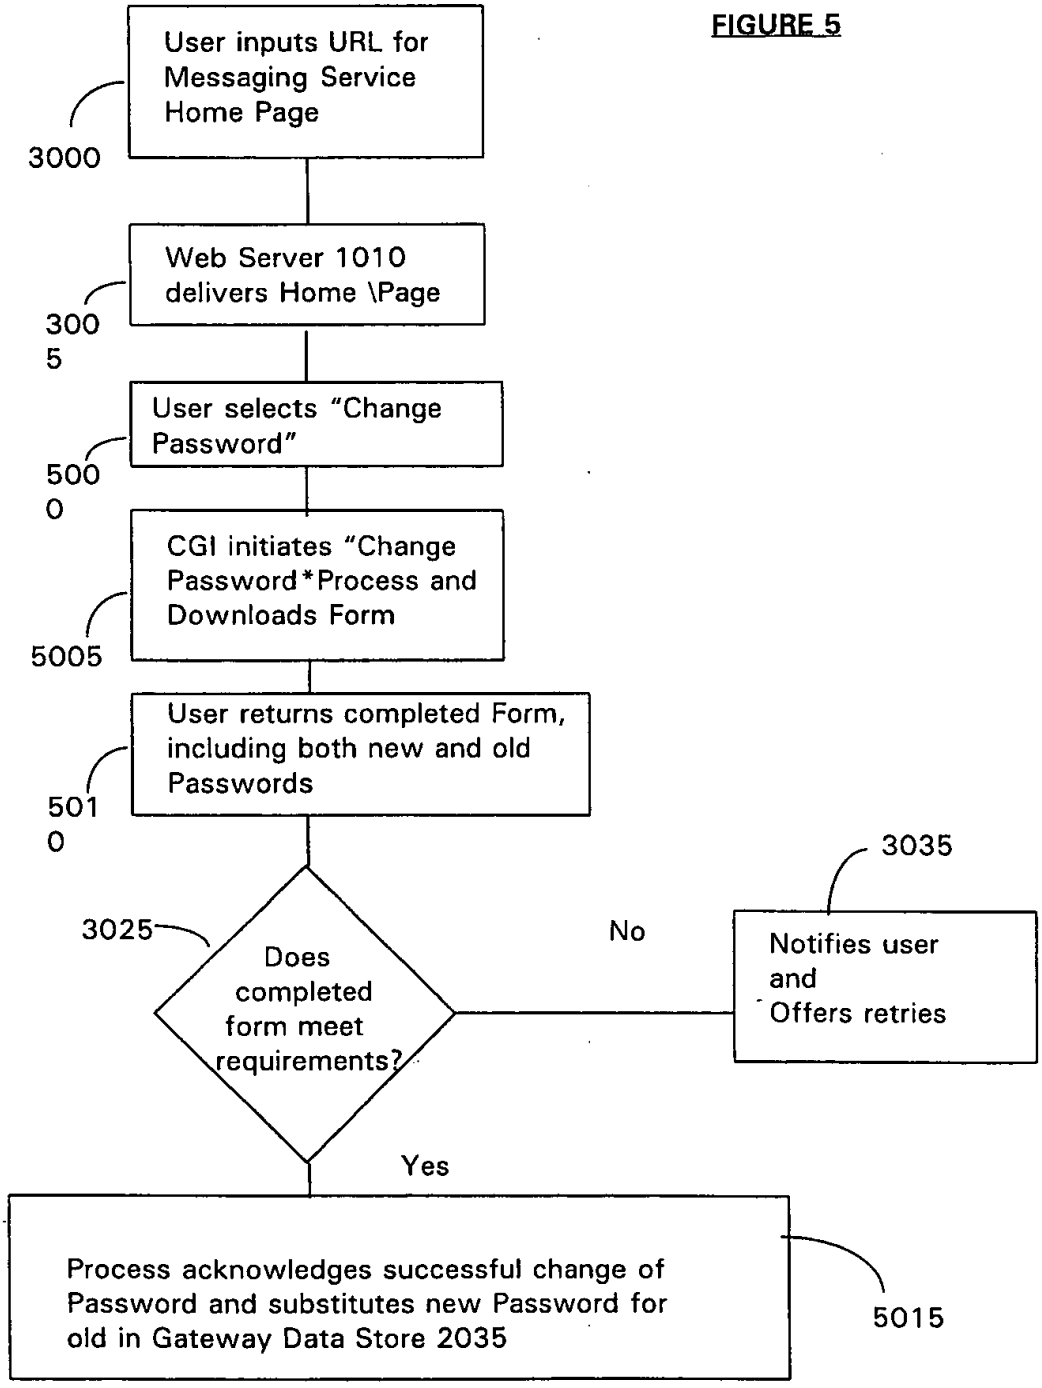
User inputs URL for Messaging Service Home Page
3000
Web Server 1010 delivers Home \Page
3005
User selects "Change Password"
5000
CGI initiates "Change Password*Process and Downloads Form
5005
User returns completed Form, including both new and old Passwords
5010
3025
Does completed form meet requirements?
No
3035
Notifies user and Offers retries
Yes
Process acknowledges successful change of Password and substitutes new Password for old in Gateway Data Store 2035
5015

FIGURE 6

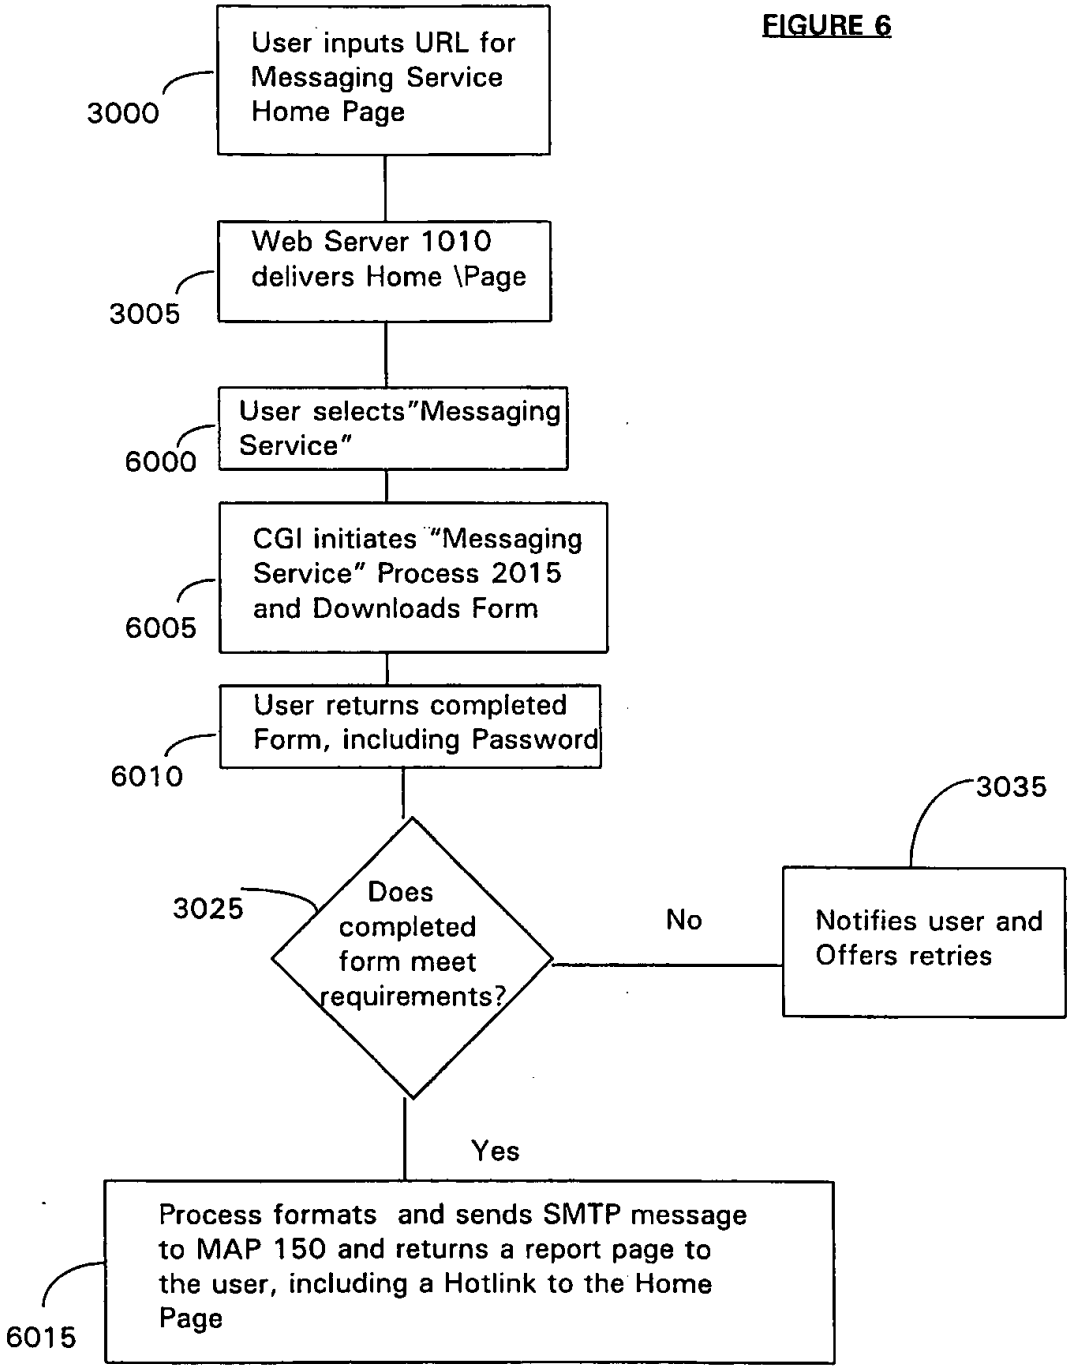
User inputs URL for Messaging Service Home Page
3000
Web Server 1010 delivers Home \Page
3005
User selects"Messaging Service"
6000
CGI initiates "Messaging Service" Process 2015 and Downloads Form
6005
User returns completed Form, including Password
6010
3035
3025
Does completed form meet requirements?
No
Notifies user and Offers retries
Yes
Process formats and sends SMTP message to MAP 150 and returns a report page to the user, including a Hotlink to the Home Page
6015

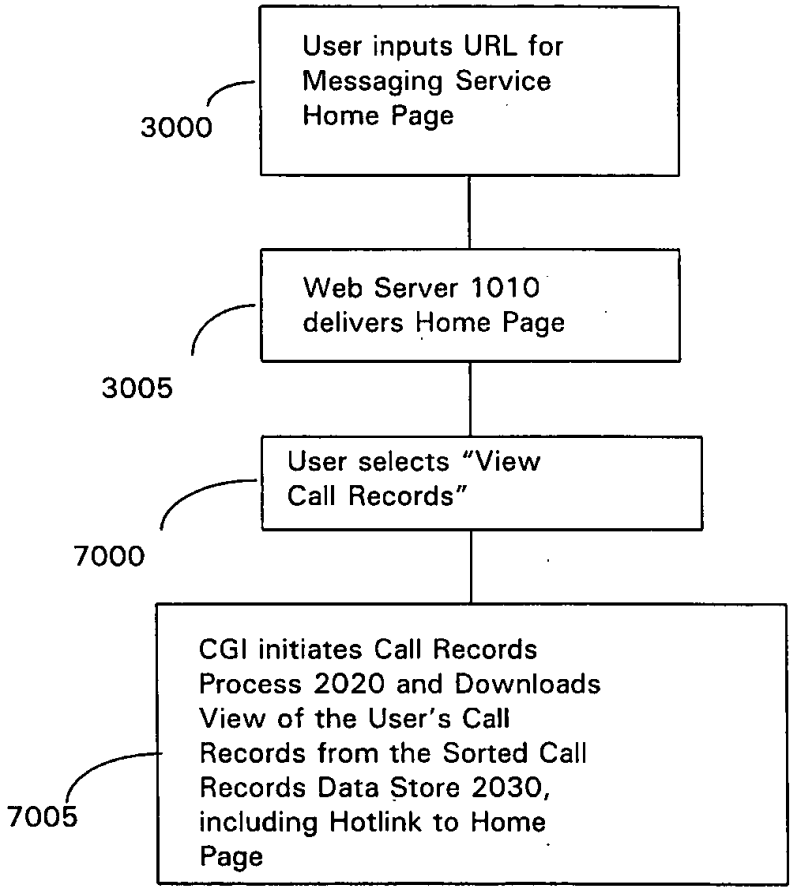
User inputs URL for Messaging Service Home Page
3000
Web Server 1010 delivers Home Page
3005
User selects "View Call Records"
7000
CGI initiates Call Records Process 2020 and Downloads View of the User's Call Records from the Sorted Call Records Data Store 2030, including Hotlink to Home Page
7005

FIGURE 7

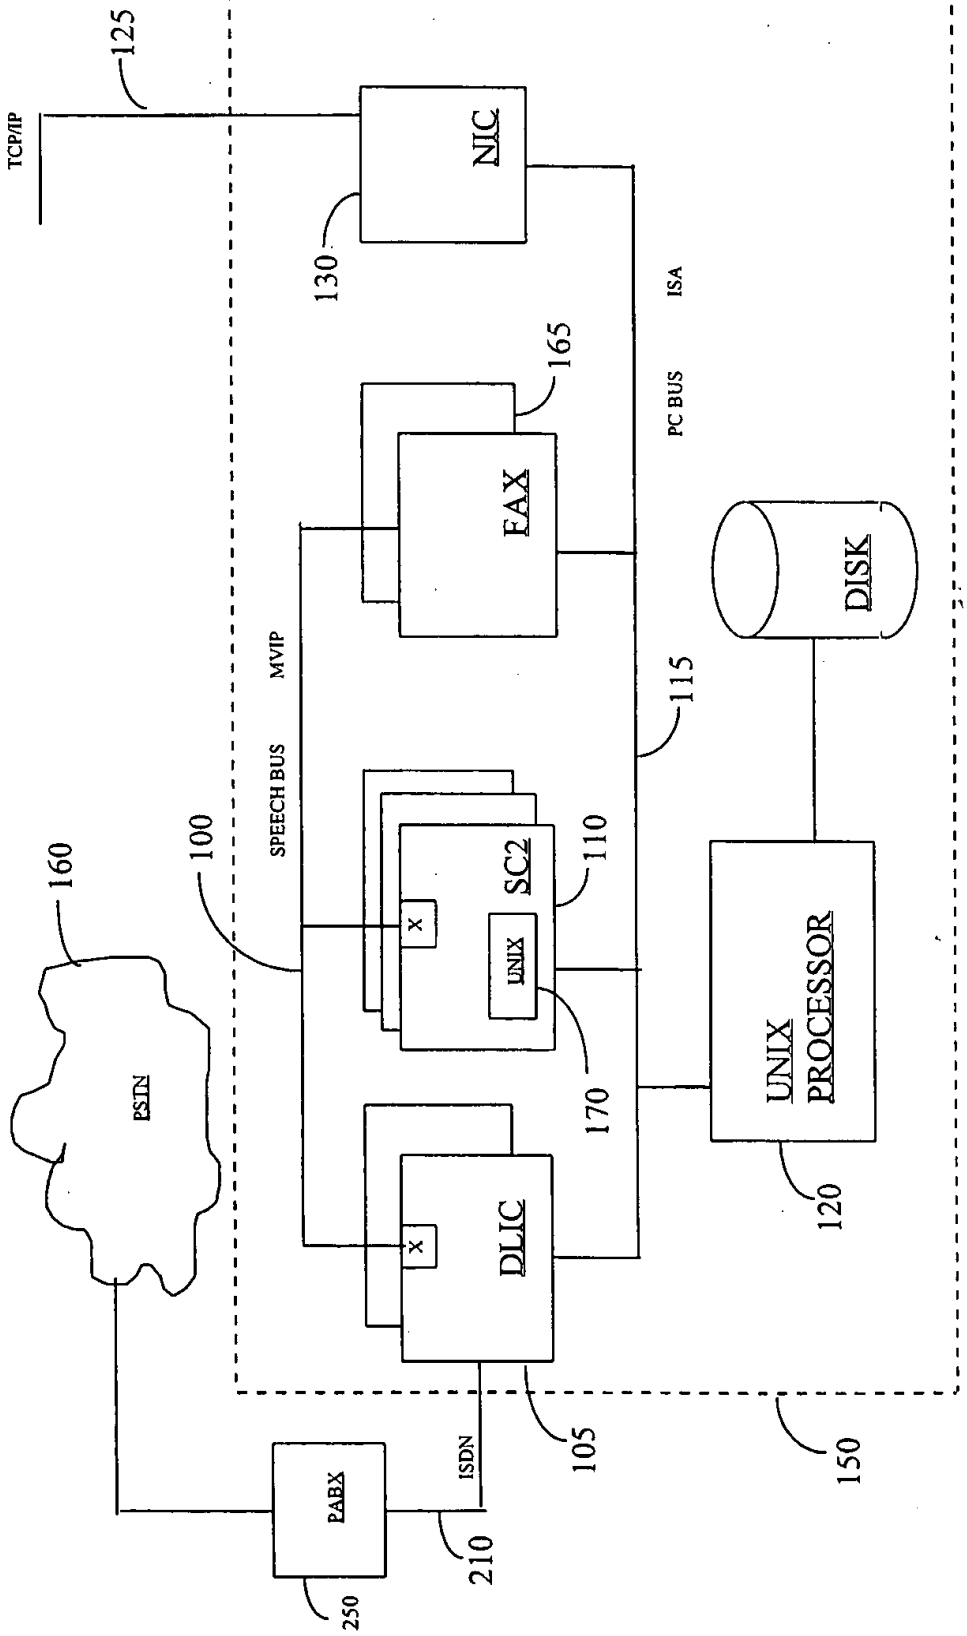
TCP/IP
125
NIC
130
ISA
PC BUS
165
FAX
DISK
115
MVIP
SPEECH BUS
100
160
SC2
110
X
UNIX
170
PSTN
UNIX
PROCESSOR
120
DLIC
105
150
ISDN
PABX
210
250

FIGURE 8

FIGURE 9

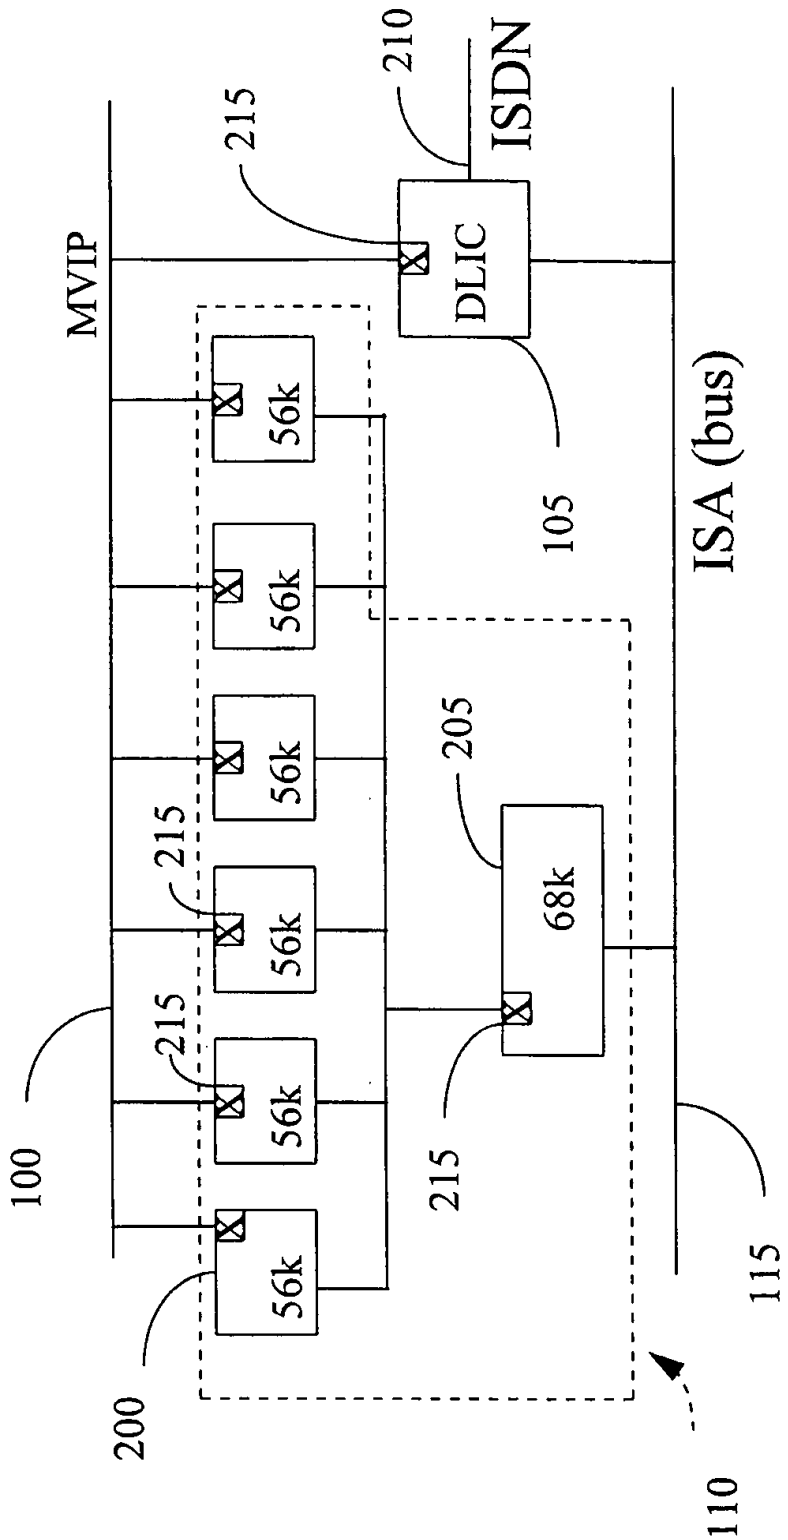

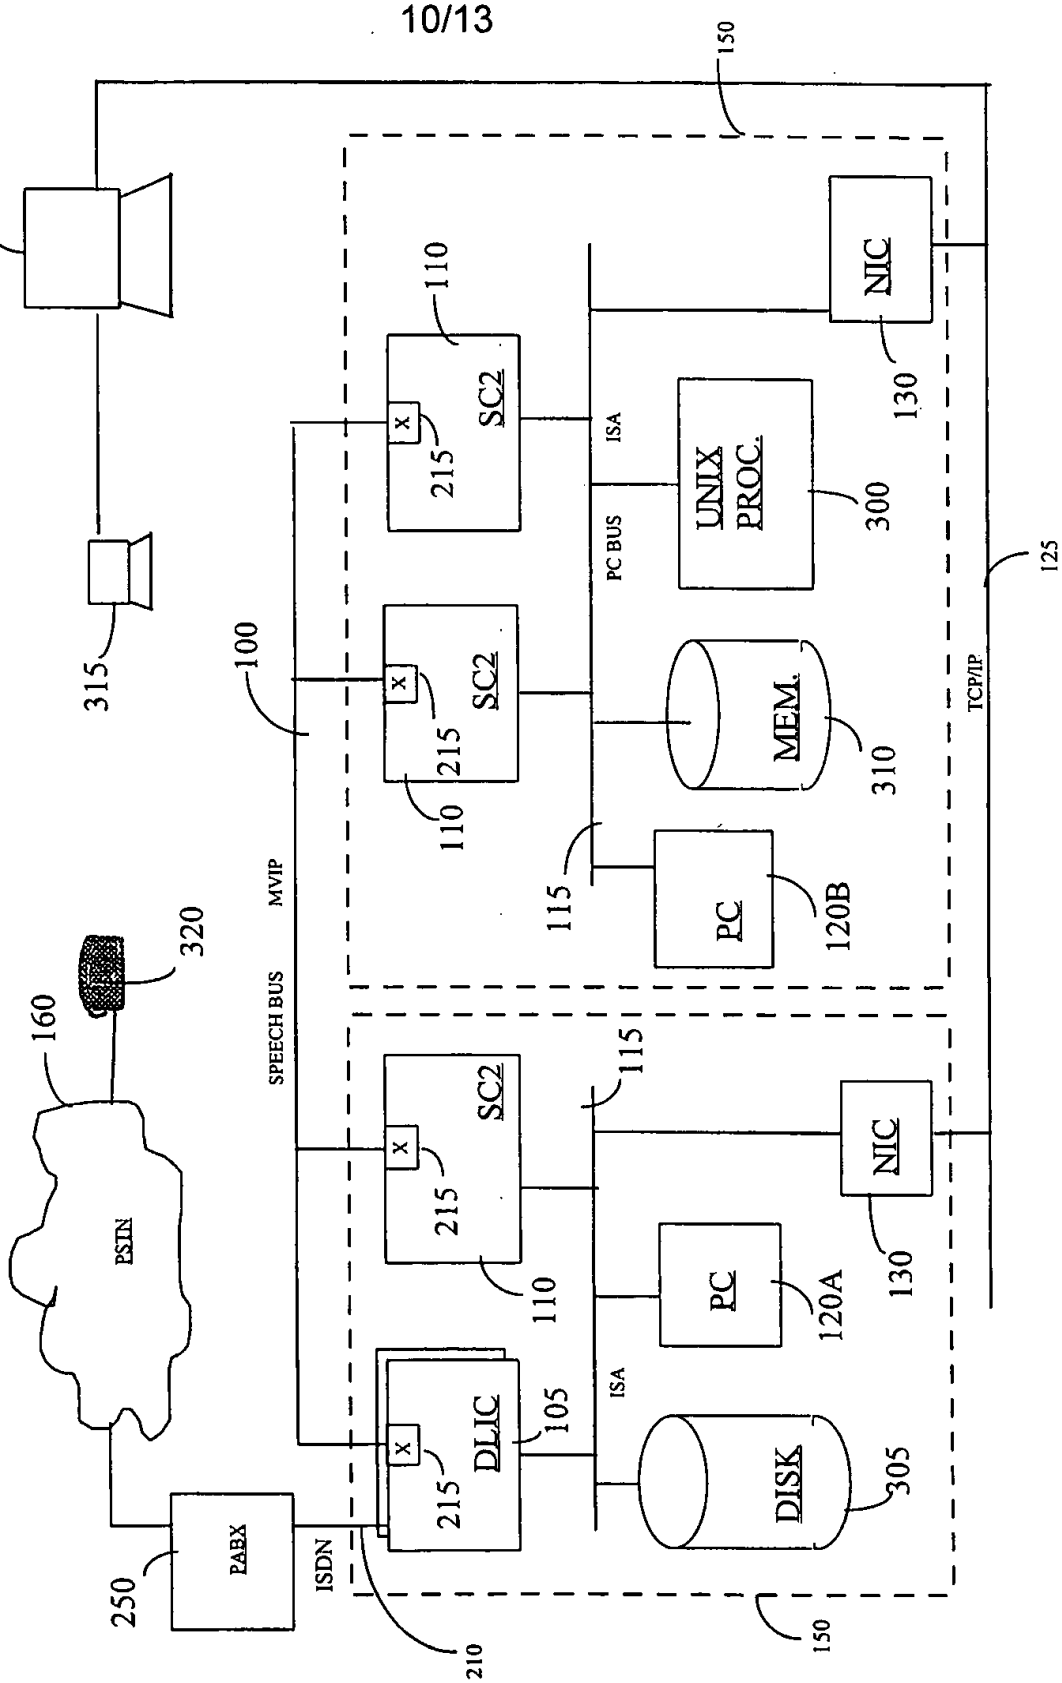
1000
150
315
100
SC2
215
110
ISA
PC BUS
UNIX PROC.
300
NIC
130
MEM.
310
TCP/IP
125
MVIP
SPEECH BUS
115
PC
120B
320
160
PSTN
DLIC
105
PC
120A
DISK
305
PABX
250
ISDN
210

FIGURE 10

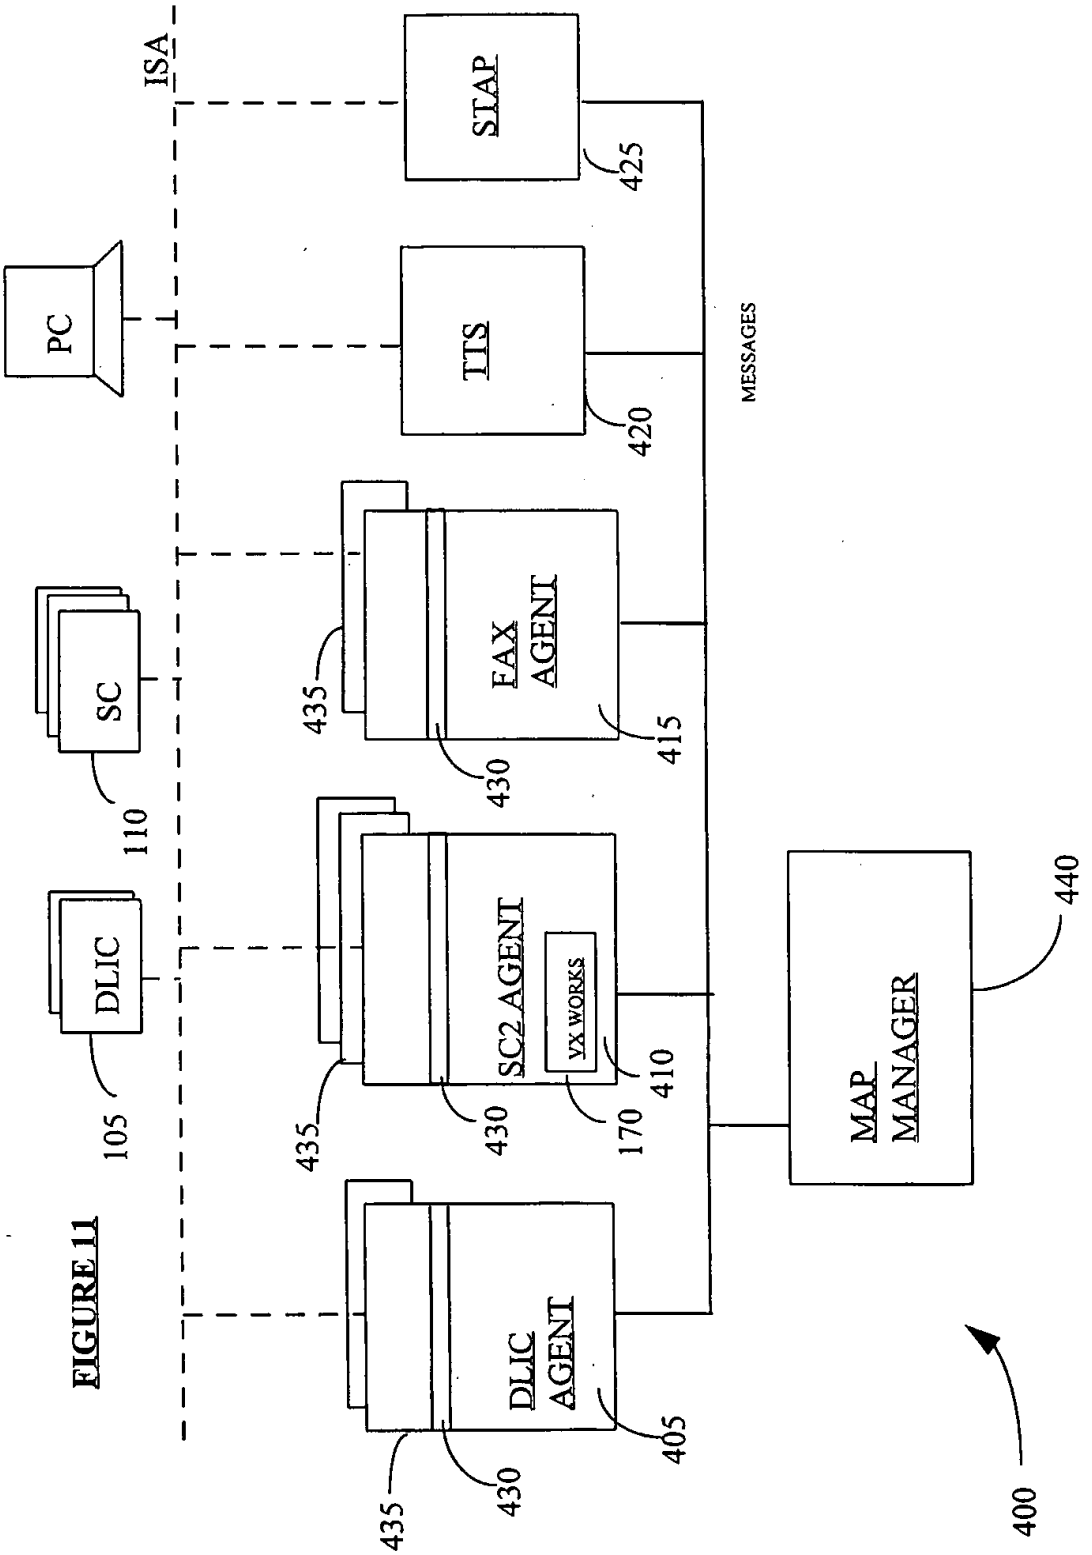
ISA
PC
SC
DLIC
105
110
STAP
425
TTS
420
MESSAGES
FAX AGENT
415
435
430
SC2 AGENT
VX WORKS
410
170
DLIC AGENT
405
MAP MANAGER
440
400

FIGURE 11

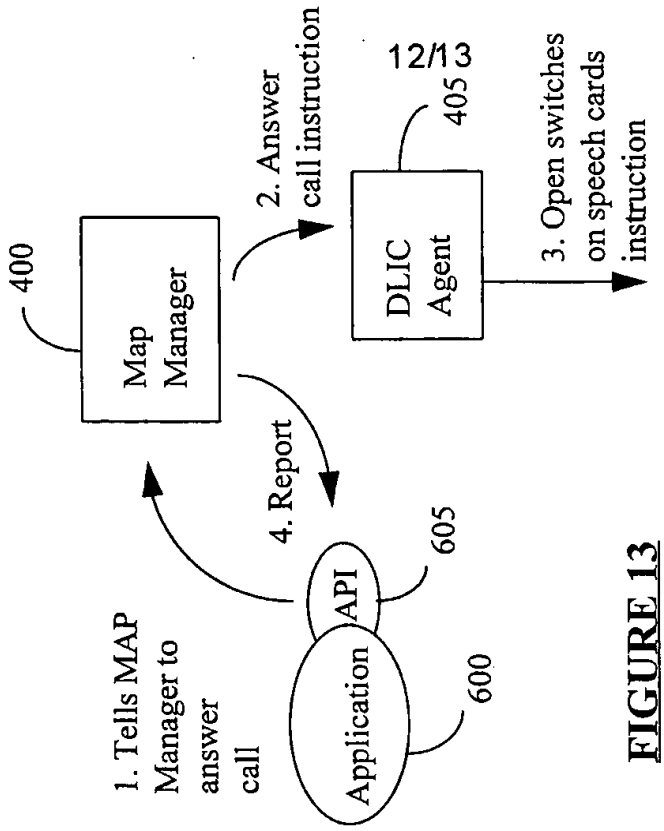

FIGURE 13

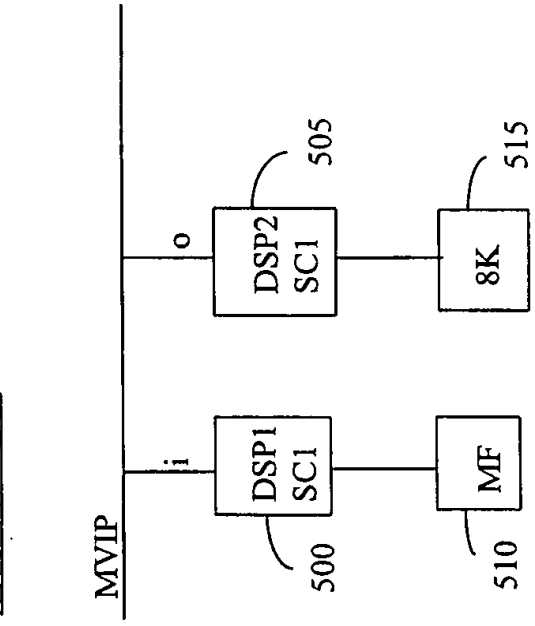

FIGURE 12

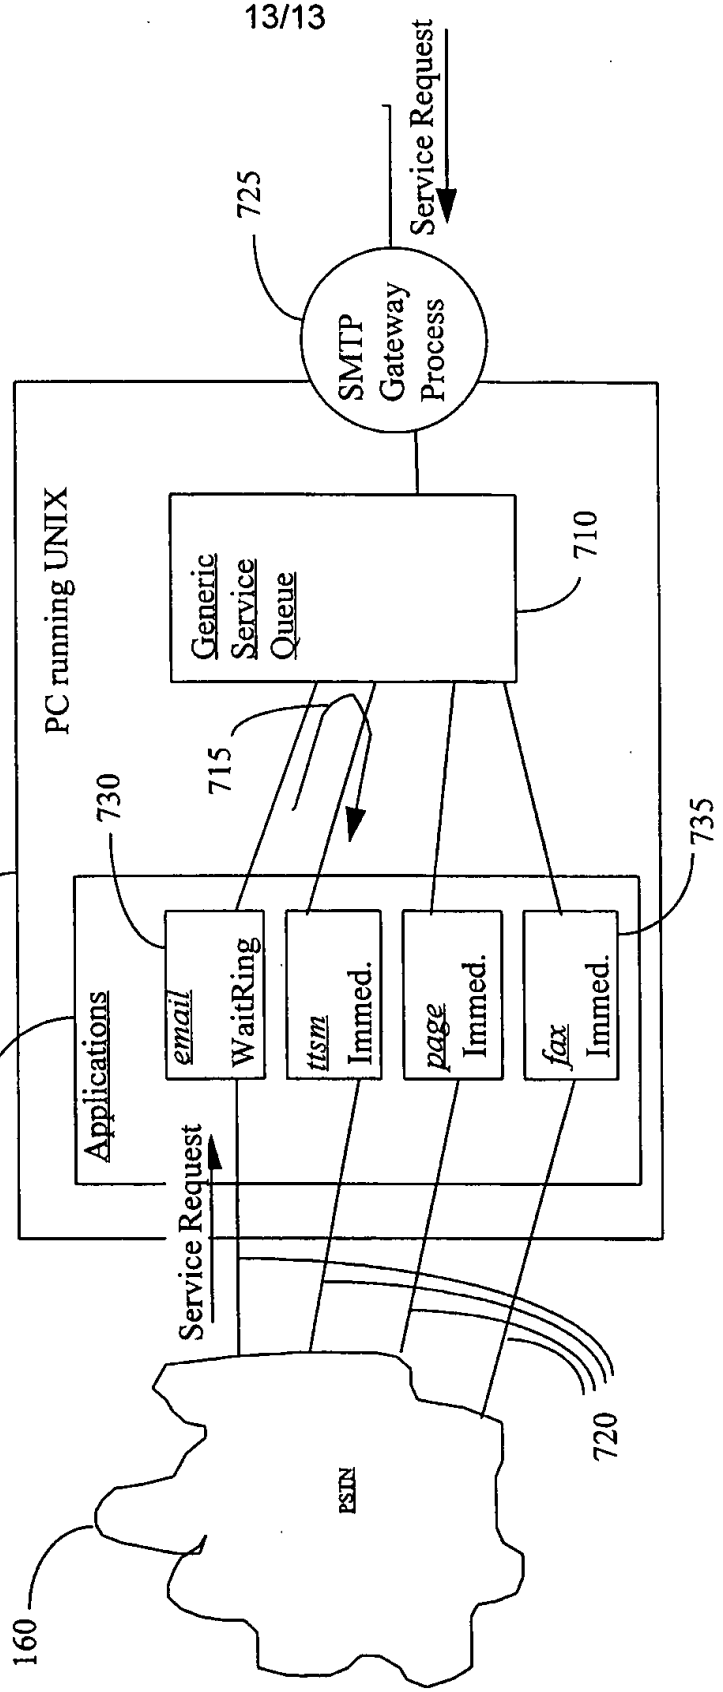
Service Request
725
SMTP Gateway Process
PC running UNIX
Generic Service Queue
710
715
730
735
400
700
Applications
email WaitRing
ttsm Immed.
page Immed.
fax Immed.
Service Request
720
PSTN
160

FIGURE 14

INTERNATIONAL SEARCH REPORT

Inter ›nal Application No
PCT/GB 98/00126

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 H04M3/50 H04M7/00 H04L12/58 H04Q3/00

According to International Patent Classification(IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 6 H04M H04L

Documentation searched other than minimumdocumentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 96 09710 A (OCTEL COMMUNICATIONS CORP) 28 March 1996 see page 1 - page 15 see page 70, line 17 - page 72, line 26 ---	1,2,4-7, 12-15, 17,22
X	WO 94 06230 A (OCTUS INC) 17 March 1994 see abstract see page 9, line 7 - line 23 see page 24, line 5 - line 19 see page 33, line 25 - page 38, line 11 --- -/--	1,2,4-6, 9-16,21, 22

[X] Further documents are listed in the continuation of box C.

[X] Patent family members are listed in annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publicationdate of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

1

Date of the actual completion of theinternational search	Date of mailing of the international search report
18 May 1998	26/05/1998
Name and mailing address of the ISA European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax: (+31-70) 340-3016	Authorized officer Megalou, M

Form PCT/ISA/210 (second sheet) (July 1992)

INTERNATIONAL SEARCH REPORT

Inter. ...nal Application No
PCT/GB 98/00126

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication,where appropriate, of the relevant passages	Relevant to claim No.
X	LOW C.: "Internet telephony red herring" HP LABORATORIES TECHNICAL REPORT, no. 96-98, June 1996, PALO ALTO, CA, USA, pages 1-15, XP002043669 see the whole document	1,14,15, 17,18
X	EP 0 631 419 A (VMX INC) 28 December 1994 see abstract; figures 7,11,13,17 see page 5, line 10 - page 6, line 12	1,5,6
P,X	WO 97 13352 A (NORTHERN TELECOM LTD) 10 April 1997 see abstract see page 16, line 27 - line 36 see page 18, line 7 - line 19 see figures 1,2	1,5,6, 13-15,22
P,X	US 5 633 916 A (ELROD EDWIN M ET AL) 27 May 1997 see abstract see claims 1,10,12	1,5,6, 12-17
P,A	US 5 610 910 A (LIVERMORE FREDERICK C ET AL) 11 March 1997 see the whole document	1-22
A	WO 95 26092 A (ERICSSON MESSAGING SYSTEMS INC) 28 September 1995 see abstract	1-22
A	WO 93 07566 A (MOTOROLA INC) 15 April 1993 see abstract	1-22
A	GB 2 301 260 A (IBM) 27 November 1996 see abstract	1-22

Form PCT/ISA/210 (continuation of second sheet) (July 1992)

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No
PCT/GB 98/00126

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 9609710 A	28-03-96	US 5740231 A	14-04-98
		US 5621727 A	15-04-97
		AU 3633795 A	09-04-96
		CA 2199243 A	28-03-96
		EP 0781482 A	02-07-97
WO 9406230 A	17-03-94	AU 5100593 A	29-03-94
EP 0631419 A	28-12-94	AT 136183 T	15-04-96
		CA 2115822 A	23-12-94
		DE 69400115 D	02-05-96
		DE 69400115 T	14-11-96
		US 5557659 A	17-09-96
		US 5717742 A	10-02-98
WO 9713352 A	10-04-97	AU 6729596 A	28-04-97
US 5633916 A	27-05-97	NONE	
US 5610910 A	11-03-97	WO 9707625 A	27-02-97
WO 9526092 A	28-09-95	US 5724406 A	03-03-98
		AU 1978595 A	09-10-95
WO 9307566 A	15-04-93	EP 0560965 A	22-09-93
		JP 6503694 T	21-04-94
GB 2301260 A	27-11-96	EP 0830779 A	25-03-98
		WO 9637989 A	28-11-96